Fig. 8. *I*-generator and summer block.

Resolver Logic Operation

Determining the number of responders following a search is basically a two-step operation sequenced by the control pulses, E_Y and $E_{Y \cdot I}$:

1) E_Y gates the Y vector out to the $A_0 - A_{192}$ inputs to the tree. A_{out} at the root of the tree then contains the value of ORing Y_i 's, y , which is temporarily stored in the combinatorial logic and store block. Simultaneously, the "0" level at the I input to the root of the tree ripples out to the $B_{0,i}$ -level block where it is combined with the $A_0 - A_{192}$ inputs to form the I vector.

2) $E_{Y \cdot I}$ gates $Y \cdot I$ over the $A_0 - A_{192}$ inputs to the tree. A_{out} at the root of the tree then contains the value of ORing $(Y \cdot I)_i$'s, which is stored and combined with the value of y as previously described, to determine the number of responders.

CONCLUSION

An optimal framing strategy and an associative processor implementation which can eliminate loss of frame synchronization due to bit errors in the framing pattern and reduce time to reframe following loss of bit count integrity to its minimum value have been described. The technique eliminates the time spent dwelling at false frame positions by simultaneously searching all possible bit positions for the frame pattern. For T1 digital terminals with frame lengths of 193 bits, a period of 125 μ s, and a 1 bit frame pattern, the reframe time can be reduced from 49.5 ms to 1.5 ms. The loss in information is thus reduced from 396 frames to 12 frames, providing a significant improvement in system performance.

ACKNOWLEDGMENT

Credit is due to G. A. Anderson for the design of the I -generator tree used for the multiple match resolution logic.

REFERENCES

- [1] Technical Staff, Bell Telephone Labs, *Transmission Systems for Communications*, 4th ed. Winston-Salem, NC: Western Electric Co., Tech. Publ., 1970, 759 pp.
- [2] L. B. W. Jolley, *Summation of Series*. New York: Dover, 1961.
- [3] G. A. Anderson, "Multiple match resolvers: A new design method," *IEEE Trans. Comput.* (Corresp.), vol. C-23, pp. 1317-1320, Dec. 1974.

The Multidimensional Access Memory in STARAN

KENNETH E. BATCHER

Abstract—STARAN® has a number of array modules, each with a multidimensional access (MDA) memory. The implementation of this memory with random-access memory (RAM) chips is de-

Manuscript received August 22, 1975; revised June 2, 1976.

The author is with the Department of Digital Technology, Goodyear Aerospace Corporation, Akron, OH 44315.

*STARAN is a registered trademark of Goodyear Aerospace Corporation, Akron, OH.

scribed. Because data can be accessed in either the word direction or the bit-slice direction, associative processing is possible without the need for costly, custom-made logic-in-memory chips.

Index Terms—Associative processing, corner-turning memories, multidimensional access (MDA) memories, parallel processing, solid-state memories.

INTRODUCTION

With the current interest in storage schemes and memory-to-processing element interconnections [1], [2], it is timely to describe the multidimensional access (MDA) memory used in the STARAN associative array processor. Conventional random-access memory (RAM) integrated-circuit chips are used in a novel manner to construct a memory that provides access to data in a number of different ways (by words, by bit slices, by bytes, etc.). The use of these standard, high-volume, low pin-count memory devices in place of custom LSI devices reduces the cost of associative processing significantly; the reason is that the ratio of the number of bits stored in a chip over the number of pins on the chip can be much higher.

NEED FOR MULTIDIMENSIONAL ACCESS

In STARAN, vector arithmetic and search operations are performed using many simple processing elements (PE's) [3]–[5]. Memory data are accessed by bit slices (one bit of many items fetched or stored in one memory cycle), with each bit assigned to a distinct PE. For example, consider the vector addition, $A + B \rightarrow B$, where each vector has 20 ten-bit components. First, the least-significant bit (LSB) of every component of A is fetched with a bit-slice access and sent to the PE's (the i th PE receiving the LSB of the i th component of A). Then the LSB of every B component is fetched with a bit-slice access and sent to the PE's (the i th PE receiving the LSB of the i th component).

In every PE, the two LSB's are added to generate a carry and the LSB of a sum component. The sum bits are sent to the LSB position of every component of B with a bit-slice store (the i th component of B receives the sum bit of the i th PE). These basic steps are repeated for each of the other nine bits of A and B , working from the low-order bit positions to the high-order bit positions. The PE's save the carries of each iteration and add them to the two operand bits of the following iteration. The 20 ten-bit additions are performed in 30 memory cycles: 10 bit-slice fetches from vector A , 10 bit-slice fetches from vector B , and 10 bit-slice stores into vector B .

Like conventional computers, input and output operations in STARAN usually access memory data by items or words (many bits of one item fetched or stored in one memory cycle) since most peripherals transmit data in this fashion. Nonvector (scalar) arithmetic operations on single items also access data by words.

Besides bit-slice accesses and word accesses, other ways of accessing memory data are useful in certain situations. For example, a numerical alphanumeric records are stored in memory. Input and output operations can access several contiguous bytes of one record in one memory cycle. Operations searching for special codes in key positions can access one byte of several stored records in one memory cycle. Operations searching for alphabetic data, numeric data, lower case data, etc., anywhere in a record can access one bit of several bytes in one memory cycle.

Thus, there is a need for a memory with a number of access modes allowing data to be accessed by bit slice, by word, etc. Such a memory allows the processing power of the system to be applied either to a few bits of many items for parallel vector operations or to many bits of a few items for conventional scalar operations.

In STARAN, MDA memories are implemented with standard RAM integrated-circuit chips—the same chips used by a number of conventional computers for solid-state storage.

MDA MEMORY CONSTRUCTION

General

The construction of an MDA memory from RAM chips can be illustrated by an example of an MDA memory with 2^n words, each of 2^m bits. The memory uses 2^n RAM chips with 2^m bit locations in each chip. Memory words, RAM chips, word bits, and bit locations in a chip are numbered (addressed) with n -bit binary vectors (n -vectors).

Notation

If X and Y are n -vectors, $X \oplus Y$ (each coordinate of X ANDed with the EXCLUSIVE OR of the corresponding coordinates of Y) is denoted $X \oplus Y$.

Storage


The MDA memory is constructed by storing to each word B the sum of the i th component of A and the i th component of B (the i th component of B receives the sum bit of the i th PE).

Fig. 1 illustrates the construction of the MDA memory. The i th component of A is stored in the i th row of the MDA memory. The i th component of B is stored in the i th column of the MDA memory. The i th component of B is stored in the i th column of the MDA memory.

Access

To access the MDA memory, the i th component of A is fetched from the i th row of the MDA memory. The i th component of B is fetched from the i th column of the MDA memory.

Every state of the MDA memory is specified by a local address. The local address is a binary vector of length n . The local address is specified by the i th component of the local address.

Fig. 8. *I*-generator and summer block.

Operation

number of responders following a search is basically sequenced by the control pulses, E_Y and $E_{Y \cdot I}$:

Y vector out to the $A_0 - A_{192}$ inputs to the tree. A_{out} tree then contains the value of ORing Y_i 's, y , which is in the combinatorial logic and store block. Simultaneously at the I input to the root of the tree ripples out to where it is combined with the $A_0 - A_{192}$ inputs to form

$Y \cdot I$ over the $A_0 - A_{192}$ inputs to the tree. A_{out} at the tree contains the value of ORing $(Y \cdot I)_i$'s, which is stored in the value of y as previously described, to determine responders.

CONCLUSION

ing strategy and an associative processor implement-
 terminate loss of frame synchronization due to bit errors
 tern and reduce time to reframe following loss of bit
 minimum value have been described. The technique
 spent dwelling at false frame positions by simulta-
 all possible bit positions for the frame pattern. For T1
 with frame lengths of 193 bits, a period of 125 μ s, and
 rn, the reframe time can be reduced from 49.5 ms to
 information is thus reduced from 396 frames to 12
 significant improvement in system performance.

ACKNOWLEDGMENT

G. A. Anderson for the design of the *I*-generator tree
 the match resolution logic.

REFERENCES

- Bell Telephone Labs, *Transmission Systems for Communi-*
 Winston-Salem, NC: Western Electric Co., Tech. Publ., 1970.
Summation of Series. New York: Dover, 1961.
 Multiple match resolvers: A new design method," *IEEE Trans.*
), vol. C-23, pp. 1317-1320, Dec. 1974.

Multidimensional Access Memory in STARAN

KENNETH E. BATCHER

STARAN® has a number of array modules, each with
 al access (MDA) memory. The implementation
 with random-access memory (RAM) chips is de-

August 22, 1975; revised June 2, 1976.

the Department of Digital Technology, Goodyear Aerospace
 OH 44315.

Registered trademark of Goodyear Aerospace Corporation, Akron,

scribed. Because data can be accessed in either the word direction
 or the bit-slice direction, associative processing is possible without
 the need for costly, custom-made logic-in-memory chips.

Index Terms—Associative processing, corner-turning memories,
 multidimensional access (MDA) memories, parallel processing,
 solid-state memories.

INTRODUCTION

With the current interest in storage schemes and memory-to-pro-
 cessing element interconnections [1], [2], it is timely to describe the
 multidimensional access (MDA) memory used in the STARAN asso-
 ciative array processor. Conventional random-access memory (RAM)
 integrated-circuit chips are used in a novel manner to construct a memory
 that provides access to data in a number of different ways (by words, by
 bit slices, by bytes, etc.). The use of these standard, high-volume, low
 pin-count memory devices in place of custom LSI devices reduces the cost
 of associative processing significantly; the reason is that the ratio of the
 number of bits stored in a chip over the number of pins on the chip can
 be much higher.

NEED FOR MULTIDIMENSIONAL ACCESS

In STARAN, vector arithmetic and search operations are performed
 using many simple processing elements (PE's) [3]-[5]. Memory data are
 accessed by bit slices (one bit of many items fetched or stored in one
 memory cycle), with each bit assigned to a distinct PE. For example,
 consider the vector addition, $A + B \rightarrow B$, where each vector has 250
 ten-bit components. First, the least-significant bit (LSB) of every
 component of A is fetched with a bit-slice access and sent to the PE's (the i th
 PE receiving the LSB of the i th component of A). Then the LSB of every
 B component is fetched with a bit-slice access and sent to the PE's (the
 i th PE receiving the LSB of the i th component).

In every PE, the two LSB's are added to generate a carry and the LSB
 of a sum component. The sum bits are sent to the LSB position of every
 component of B with a bit-slice store (the i th component of B receives
 the sum bit of the i th PE). These basic steps are repeated for each of the
 other nine bits of A and B , working from the low-order bit positions to
 the high-order bit positions. The PE's save the carries of each iteration
 and add them to the two operand bits of the following iteration. The 250
 ten-bit additions are performed in 30 memory cycles: 10 bit-slice fetches
 from vector A , 10 bit-slice fetches from vector B , and 10 bit-slice stores
 into vector B .

Like conventional computers, input and output operations in STAR-
 AN usually access memory data by items or words (many bits of one item
 fetched or stored in one memory cycle) since most peripherals transmit
 data in this fashion. Nonvector (scalar) arithmetic operations on single
 items also access data by words.

Besides bit-slice accesses and word accesses, other ways of accessing
 memory data are useful in certain situations. For example, a number of
 alphanumeric records are stored in memory. Input and output operations
 can access several contiguous bytes of one record in one memory cycle.
 Operations searching for special codes in key positions can access one byte
 of several stored records in one memory cycle. Operations searching for
 alphabetic data, numeric data, lower case data, etc., anywhere in a record,
 can access one bit of several bytes in one memory cycle.

Thus, there is a need for a memory with a number of access modes al-
 lowing data to be accessed by bit slice, by word, etc. Such a memory allows
 the processing power of the system to be applied either to a few bits of
 many items for parallel vector operations or to many bits of a few items
 for conventional scalar operations.

In STARAN, MDA memories are implemented with standard RAM
 integrated-circuit chips—the same chips used by a number of conven-
 tional computers for solid-state storage.

MDA MEMORY CONSTRUCTION

General

The construction of an MDA memory from RAM chips can be illus-
 trated by an example of an MDA memory with 2^n words, each of 2^m bits.
 The memory uses 2^n RAM chips with 2^m bit locations in each chip.
 Memory words, RAM chips, word bits, and bit locations in a chip are
 numbered (addressed) with n -bit binary vectors (n -vectors).

Notation

If X and Y are two n -vectors, then \bar{X}
 X (each component bit is replaced by
 cates the logical product (AND) of X
 ANDed with the corresponding compo
 the EXCLUSIVE-OR of X and Y (a c
 if the corresponding components of X

Storage Pattern

The MDA memory stores data in a
 according to the following rule.

Storage Rule: For any two n -vector
 stored in bit-location B of memory c
 two- n -vectors, B and C , bit-location B
 word $B \oplus C$.

Fig. 1 illustrates the storage rule for
 i is stored in bit-location i of all eight
 0 and 7 occupy the main forward and
 array, and the other six words occupy th
 of certain subarrays. Each bit-slice an
 eight memory chips. This scrambling
 various ways.

Access Stencils

To access data in the MDA memory,
 stencil is specified. A word stencil cov
 positioned over any memory word. A
 (one bit of each word) and can be po
 stencils cover some bits of some words
 bits of certain words.

Every stencil covers exactly 2^n bits
 fetched or stored in one memory cycle
 In general, there are 2^n possible ste
 positioned in 2^n different places. Th
 n -vector called the access mode, and
 vector called the global address. When
 are specified, the address bus structu
 a local address for each memory chip

Local Addressing Rule: For any th
 the specified access mode and G is the
 local address fed to the address pins c

er the word direction
ng is possible without
emory chips.

er-turning memories,
parallel processing,

s and memory-to-pro-
timely to describe the
in the STARAN asso-
ccess memory (RAM)
t to construct a memory
ent ways (by words, by
ard, high-volume, low
devices reduces the cost
is that the ratio of the
of pins on the chip can

ACCESS

erations are performed
-[5]. Memory data are
ched or stored in one
inct PE. For example,
e each vector has 250
it (LSB) of every com-
ent to the PE's (the i th
Then the LSB of every
d sent to the PE's (the

te a carry and the LSB
LSB position of every
omponent of B receives
epeated for each of the
-order bit positions to
arries of each iteration
ving iteration. The 250
les; 10 bit-slice fetches
and 10 bit-slice stores

operations in STAR-
(many bits of one item
peripherals transmit
operations on single

her ways of accessing
example, a number of
and output operations
in one memory cycle.
ns can access one byte
erations searching for
anywhere in a record,
cycle.

er of access modes al-
Such a memory allows
either to a few bits of
ny bits of a few items

with standard RAM
a number of conven-

ION

M chips can be illus-
ords, each of 2^n bits.
ations in each chip.
ations in a chip are
-vectors).

0	a_{00}	a_{11}	a_{22}	a_{33}	a_{44}	a_{55}	a_{66}	a_{77}
1	a_{01}	a_{10}	a_{23}	a_{32}	a_{45}	a_{54}	a_{67}	a_{76}
2	a_{02}	a_{13}	a_{20}	a_{31}	a_{46}	a_{57}	a_{64}	a_{75}
3	a_{03}	a_{12}	a_{21}	a_{30}	a_{47}	a_{56}	a_{65}	a_{74}
4	a_{04}	a_{15}	a_{26}	a_{37}	a_{40}	a_{51}	a_{62}	a_{73}
5	a_{05}	a_{14}	a_{27}	a_{36}	a_{41}	a_{50}	a_{63}	a_{72}
6	a_{06}	a_{17}	a_{24}	a_{35}	a_{42}	a_{53}	a_{60}	a_{71}
7	a_{07}	a_{16}	a_{25}	a_{34}	a_{43}	a_{52}	a_{61}	a_{70}
	0	1	2	3	4	5	6	7

BIT-LOCATION

a_{ij} = BIT i OF WORD j

Fig. 1. Storage rule for an 8 x 8 memory.

Notation

If X and Y are two n -vectors, then \bar{X} indicates the logical negation of X (each component bit is replaced by its one's complement), XY indicates the logical product (AND) of X and Y (each component of X is ANDed with the corresponding component of Y), and $X \oplus Y$ indicates the EXCLUSIVE-OR of X and Y (a component of $X \oplus Y$ is 1 if and only if the corresponding components of X and Y have opposite values).

Storage Pattern

The MDA memory stores data in a particular scrambled pattern according to the following rule.

Storage Rule: For any two n -vectors, B and W , bit B of word W is stored in bit-location B of memory chip $B \oplus W$. Conversely, for any two n -vectors, B and C , bit-location B of memory chip C stores bit B of word $B \oplus C$.

Fig. 1 illustrates the storage rule for an 8 x 8 MDA memory. Bit-slice i is stored in bit-location i of all eight memory chips ($0 \leq i \leq 7$). Words 0 and 7 occupy the main forward and backward diagonals of the square array, and the other six words occupy the forward and backward diagonals of certain subarrays. Each bit-slice and each word are spread across the eight memory chips. This scrambling allows the data to be accessed in various ways.

Access Stencils

To access data in the MDA memory, the shape and position of an access stencil is specified. A word stencil covers all bits of one word and can be positioned over any memory word. A bit-slice stencil covers a bit-slice (one bit of each word) and can be positioned over any bit slice. Other stencils cover some bits of some words and can be positioned over certain bits of certain words.

Every stencil covers exactly 2^n bits of storage. All covered bits can be fetched or stored in one memory cycle.

In general, there are 2^n possible stencil shapes, and each shape can be positioned in 2^n different places. The stencil shape is specified by an n -vector called the access mode, and its position is specified by an n -vector called the global address. When an access mode and global address are specified, the address bus structure of the MDA memory generates a local address for each memory chip according to the following rule.

Local Addressing Rule: For any three n -vectors, M , G and C , if M is the specified access mode and G is the specified global address, then the local address fed to the address pins of memory chip C is $G \oplus (MC)$.

From the storage rule, bit-location $G \oplus (MC)$ of memory chip C contains bit $G \oplus (MC)$ of memory word $G \oplus (MC) \oplus C = G \oplus (\bar{MC})$.

Let g_k be the k th bit of global address vector G , and m_k be the k th bit of access mode vector M .

The address bus structure has $2n$ address lines (as opposed to n address lines in the conventional RAM). The address lines are labelled $x_1, y_1, \dots, x_n, y_n$. Address line x_k is driven by g_k , and address line y_k is driven by $g_k \oplus m_k$. Address pin k of memory chip $C = (c_1 c_2 \dots c_n)$ is connected either to x_k if $c_k = 0$ or to y_k if $c_k = 1$. Address line x_k feeds half of the memory chips, while y_k feeds the remaining half.

This structure implements the local addressing rule since for $k = 1, 2, \dots, n$, then $g_k \oplus m_k c_k = g_k$ if $c_k = 0$, and $g_k \oplus m_k c_k = g_k \oplus m_k$ if $c_k = 1$. All 2^n local addresses required by the memory chips are computed from just n EXCLUSIVE-OR logic circuits. Fig. 2 illustrates the address bus structure for an 8 x 8 MDA memory.

Scrambling and Unscrambling

Because of the scrambled pattern of stored data, the data fetched from the MDA memory are scrambled and, hence, must be unscrambled before being read by the processing elements (or the input/output channels). Similarly, data to be stored in memory must be scrambled before being fed to the data-input pins of the RAM chips. The processing elements and the input lines, and the output lines are numbered with n -vectors.

Scramble/Unscramble Rule: For any two n -vectors, G and P , when storing data into memory under global address G (and any access mode), the data bit on the data-input pin of processing element P (or input line P) is applied to the data-input pin of memory chip $G \oplus P$. Conversely, for any two n -vectors, G and C , when fetching memory data under global address G (and any access mode), the data bit on the data-output pin of memory chip C is sent to processing element $G \oplus C$ (or output line $G \oplus C$).

Because of the similarity between the scramble and unscramble rule and the fact that memory fetches and memory stores never occur simultaneously, one scramble/unscramble network suffices for both operations. Fig. 3 shows how the scramble/unscramble network fits between the MDA memory elements and the processing elements.

The scramble/unscramble network can be constructed using several levels of data-selector integrated circuits. With two-input selectors, there are n levels of logic with 2^n selectors in each level. The first level treats bit g_n of the global address G ; if the selectors are numbered with n -vectors, then each selector S , selects either input-data-bit S if $g_n = 0$ or input-data-bit $S \oplus (00 \dots 001)$ if $g_n = 1$. The second level of selectors treats global address bit g_{n-1} ; selector S in this level either selects first-level output S if $g_{n-1} = 0$ or selects first-level output $S \oplus (0 \dots 0010)$.

0	a_{00}	a_{11}	a_{22}	a_{33}	a_{44}	a_{55}	a_{66}	a_{77}
1	a_{01}	a_{10}	a_{23}	a_{32}	a_{45}	a_{54}	a_{67}	a_{76}
2	a_{02}	a_{13}	a_{20}	a_{31}	a_{46}	a_{57}	a_{64}	a_{75}
3	a_{03}	a_{12}	a_{21}	a_{30}	a_{47}	a_{56}	a_{65}	a_{74}
4	a_{04}	a_{15}	a_{26}	a_{37}	a_{40}	a_{51}	a_{62}	a_{73}
5	a_{05}	a_{14}	a_{27}	a_{36}	a_{41}	a_{50}	a_{63}	a_{72}
6	a_{06}	a_{17}	a_{24}	a_{35}	a_{42}	a_{53}	a_{60}	a_{71}
7	a_{07}	a_{16}	a_{25}	a_{34}	a_{43}	a_{52}	a_{61}	a_{70}
	0	1	2	3	4	5	6	7

BIT-LOCATION

a_{ij} = BIT i OF WORD j

Fig. 1. Storage rule for an 8 x 8 memory.

For any two n -vectors, X and Y , \bar{X} indicates the logical negation of X (each component of X is replaced by its one's complement), XY indicates the component-wise AND product of X and Y , and $X \oplus Y$ indicates the component-wise EXCLUSIVE-OR of X and Y (a component of $X \oplus Y$ is 1 if and only if the corresponding components of X and Y have opposite values).

Storage Rule: For any two n -vectors, B and W , bit B of word W is stored in bit-location B of memory chip $B \oplus W$. Conversely, for any two n -vectors, B and C , bit-location B of memory chip C stores bit B of word $C \oplus B$.

Access Rule: For any two n -vectors, G and P , when storing data into memory under global address G (and any access mode), the data bit from processing element P (or input line P) is applied to the data-input pin of memory chip $G \oplus P$. Conversely, for any two n -vectors, G and C , when fetching memory data under global address G (and any access mode), the data bit on the data-output pin of memory chip C is sent to processing element $G \oplus C$ (or output line $G \oplus C$).

Scrambling and Unscrambling: Because of the scrambled pattern of stored data, the data fetched from the MDA memory are scrambled and, hence, must be unscrambled before being read by the processing elements (or the input/output channels). Similarly, data to be stored in memory must be scrambled before being fed to the data-input pins of the RAM chips. The processing elements, the input lines, and the output lines are numbered with n -vectors.

Scramble/Unscramble Rule: For any two n -vectors, G and P , when storing data into memory under global address G (and any access mode), the data bit from processing element P (or input line P) is applied to the data-input pin of memory chip $G \oplus P$. Conversely, for any two n -vectors, G and C , when fetching memory data under global address G (and any access mode), the data bit on the data-output pin of memory chip C is sent to processing element $G \oplus C$ (or output line $G \oplus C$).

Access Rule: For any two n -vectors, G and P , when storing data into memory under global address G (and any access mode), the data bit from processing element P (or input line P) is applied to the data-input pin of memory chip $G \oplus P$. Conversely, for any two n -vectors, G and C , when fetching memory data under global address G (and any access mode), the data bit on the data-output pin of memory chip C is sent to processing element $G \oplus C$ (or output line $G \oplus C$).

From the storage rule, bit-location $G \oplus (MC)$ of memory chip C contains bit $G \oplus (MC)$ of memory word $G \oplus (MC) \oplus C = G \oplus (\bar{M}C)$.

Let g_k be the k th bit of global address vector G , and m_k be the k th bit of access mode vector M .

The address bus structure has $2n$ address lines (as opposed to n address lines in the conventional RAM). The address lines are labelled $x_1, y_1, x_2, \dots, x_n, y_n$. Address line x_k is driven by g_k , and address line y_k is driven by $g_k \oplus m_k$. Address pin k of memory chip $C = (c_1 c_2 \dots c_n)$ is connected either to x_k if $c_k = 0$ or to y_k if $c_k = 1$. Address line x_k feeds half of the memory chips, while y_k feeds the remaining half.

This structure implements the local addressing rule since for $k = 1, 2, \dots, n$, then $g_k \oplus m_k c_k = g_k$ if $c_k = 0$, and $g_k \oplus m_k c_k = g_k \oplus m_k$ if $c_k = 1$. All 2^n local addresses required by the memory chips are computed from just n EXCLUSIVE-OR logic circuits. Fig. 2 illustrates the address bus structure for an 8 x 8 MDA memory.

Scrambling and Unscrambling

Because of the scrambled pattern of stored data, the data fetched from the MDA memory are scrambled and, hence, must be unscrambled before being read by the processing elements (or the input/output channels). Similarly, data to be stored in memory must be scrambled before being fed to the data-input pins of the RAM chips. The processing elements, the input lines, and the output lines are numbered with n -vectors.

Scramble/Unscramble Rule: For any two n -vectors, G and P , when storing data into memory under global address G (and any access mode), the data bit from processing element P (or input line P) is applied to the data-input pin of memory chip $G \oplus P$. Conversely, for any two n -vectors, G and C , when fetching memory data under global address G (and any access mode), the data bit on the data-output pin of memory chip C is sent to processing element $G \oplus C$ (or output line $G \oplus C$).

Because of the similarity between the scramble and unscramble rules and the fact that memory fetches and memory stores never occur simultaneously, one scramble/unscramble network suffices for both operations. Fig. 3 shows how the scramble/unscramble network fits between the MDA memory elements and the processing elements.

The scramble/unscramble network can be constructed using several levels of data-selector integrated circuits. With two-input selectors, there are n levels of logic with 2^n selectors in each level. The first level treats bit g_n of the global address G ; if the selectors are numbered with n -vectors, then each selector S , selects either input-data-bit S if $g_n = 0$ or input-data-bit $S \oplus (00 \dots 001)$ if $g_n = 1$. The second level of selectors treats global address bit g_{n-1} ; selector S in this level either selects first-level output S if $g_{n-1} = 0$ or selects first-level output $S \oplus (0 \dots 0010)$ if

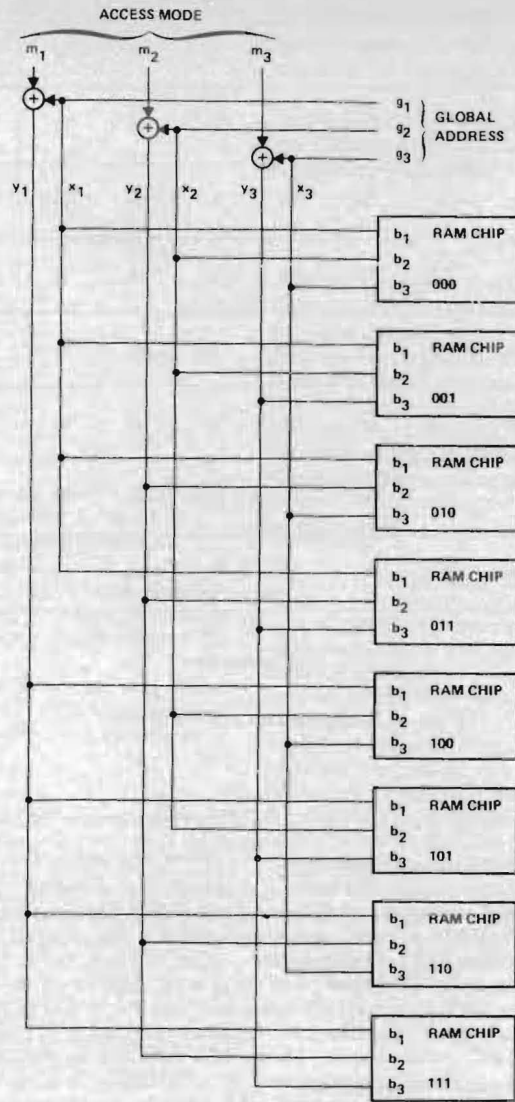


Fig. 2. Address bus structure of 8 x 8 MDA memory.

$g_{n-1} = 1$. Other levels of selectors treat the other bits of the global address in a similar manner. If four-input selectors are used, only $n/2$ levels of selectors are required; each level treats two bits of the global address.

Alternatively, the scramble/unscramble network can be constructed with one level of selection and a perfect shuffle [6]. First, the data are shuffled. Then if g_1 (the first bit of the global address) equals 1, the data in each pair of locations are swapped (data bit S goes to location $S \oplus (00 \dots 001)$). Next the data are shuffled again and, if $g_2 = 1$, then data in each pair of locations are swapped. These steps are executed n times—once for each bit of the global address.

ACCESSING THE MDA MEMORY

General

From the previously stated storage rule, local-addressing rule, and scramble/unscramble rule, one can develop an access rule that shows the memory bit accessed by any particular PE (or input-output line) when the memory is accessed with a given access mode and global address.

Access Rule: For any three n -vectors, M , G , and P , when fetching or storing data under access mode M and global address G , then processing element P (or input-output line P) will access bit $(MG) \oplus (MP)$ of memory word $(MG) \oplus (MP)$.

For bit-slice access, the access mode M equals $(00 \dots 00)$, and the global address G specifies which bit slice is fetched or stored. Processing element P accesses bit G of word P (see Fig. 4).

For word access, the access mode M equals $(11 \dots 11)$, and the global

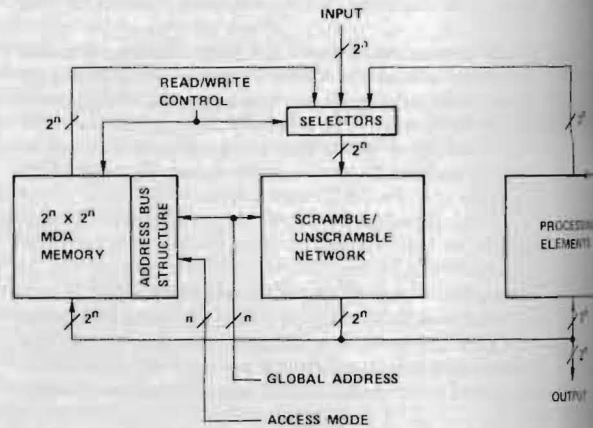


Fig. 3. Scramble/unscramble network connecting MDA memory to processing elements.

address G specifies which word is fetched or stored. Processing element P now accesses bit P of word G .

If the access mode M is $(00 \dots 00111)$ and the global address is $(b_1 b_2 \dots b_{n-3} w_{n-2} w_{n-1} w_n)$, then processing element $(p_1 p_2 \dots p_n)$ accesses bit $(b_1 b_2 \dots b_{n-3} p_{n-2} p_{n-1} p_n)$ of word $(p_1 p_2 \dots p_{n-3} w_{n-2} w_{n-1} w_n)$. One 8-bit byte from every eighth word is accessed. The first part of the global address $(b_1 b_2 \dots b_{n-3})$ specifies which

Fig. 4. MD

WORD 100 ... $00w_n 2^n$
WORD 101 ... $01w_n 2^n$
WORD 102 ... $10w_n 2^n$

WORD 111 ... $11w_n 2^n$

GLOBAL ADDR

P

2^{n-1} bytes in a word is accessed. The bit $(w_{n-3} w_{n-2} w_{n-1} w_n)$ specifies the byte-access stencil. No word is stored in memory with each bit of the global address then the byte-access stencil is used.

In general, if the access mode M covers 2^m bits of each of 2^{n-m} words of the access mode, then the 2^{n-m} covered words are

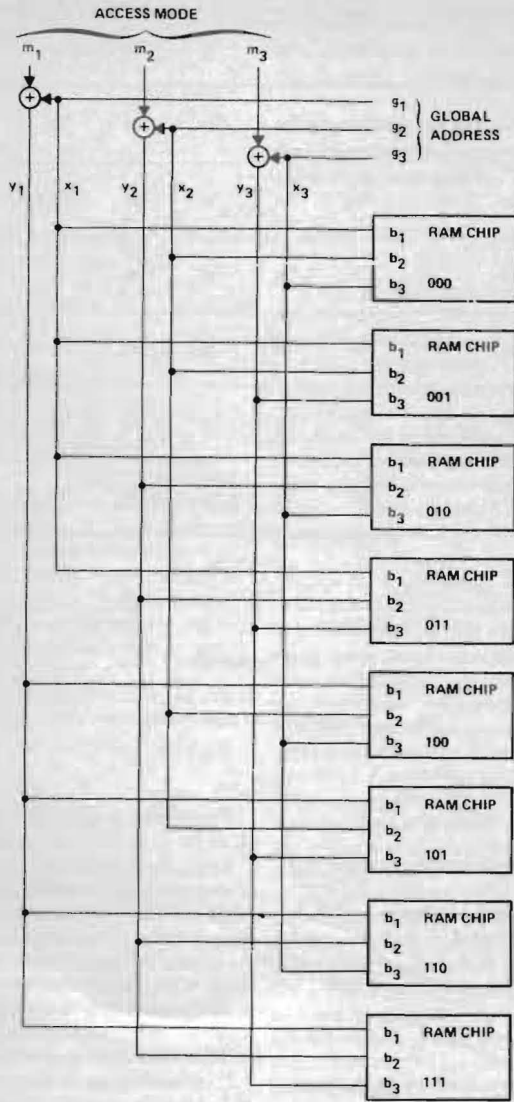


Fig. 2. Address bus structure of 8 x 8 MDA memory.

levels of selectors treat the other bits of the global address. If four-input selectors are used, only $n/2$ levels of required; each level treats two bits of the global address. The scramble/unscramble network can be constructed by selection and a perfect shuffle [6]. First, the data are swapped (data bit S goes to location $S \oplus 1$ at the data are shuffled again and, if $g_2 = 1$, then data locations are swapped. These steps are executed n times of the global address.

ACCESSING THE MDA MEMORY

Previously stated storage rule, local-addressing rule, and scramble rule, one can develop an access rule that shows the bit slice accessed by any particular PE (or input-output line) when accessed with a given access mode and global address. For any three n -vectors, M , G , and P , when fetching or storing with access mode M and global address G , then processing element P will access bit $(MG) \oplus (MP)$ of word G . If the access mode M equals $(00 \dots 00)$, and the global address G specifies which bit slice is fetched or stored. Processing element P will access bit (MP) of word G . If the access mode M equals $(11 \dots 11)$, and the global



Fig. 4. MDA memory structure.

WORD $(00 \dots 00w_{n-2}w_{n-1}w_n)$
 WORD $(00 \dots 01w_{n-2}w_{n-1}w_n)$
 WORD $(00 \dots 10w_{n-2}w_{n-1}w_n)$

WORD $(11 \dots 11w_{n-2}w_{n-1}w_n)$

GLOBAL ADDRESS = $(b_1b_2b_3)$

Fig. 5. Bit slice access.

2^{n-3} bytes in a word is accessed; the bit slice $(w_{n-2}w_{n-1}w_n)$ specifies the subset of words accessed by this byte-access stencil. Note that the data is stored in memory with each record then the byte-access stencil accesses the record.

In general, if the access mode has m bits, it covers 2^m bits of each of 2^{n-m} memory words. At the end of the access mode, the 2^m covered words are separated into

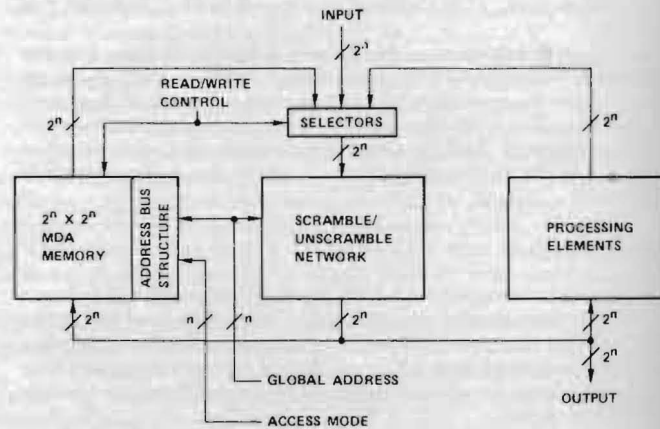


Fig. 3. Scramble/unscramble network connecting MDA memory to processing elements.

address G specifies which word is fetched or stored. Processing element P now accesses bit P of word G .

If the access mode M is $(00 \dots 00111)$ and the global address G is $(b_1b_2 \dots b_{n-3}w_{n-2}w_{n-1}w_n)$, then processing element $(p_1p_2 \dots p_{n-1}p_n)$ accesses bit $(b_1b_2 \dots b_{n-3}p_{n-2}p_{n-1}p_n)$ of word $(p_1p_2 \dots p_{n-3}w_{n-2}w_{n-1}w_n)$. One 8-bit byte from every eighth word is accessed. The first part of the global address $(b_1b_2 \dots b_{n-3})$ specifies which of the




Fig. 4. MDA memory with two-access stencils.

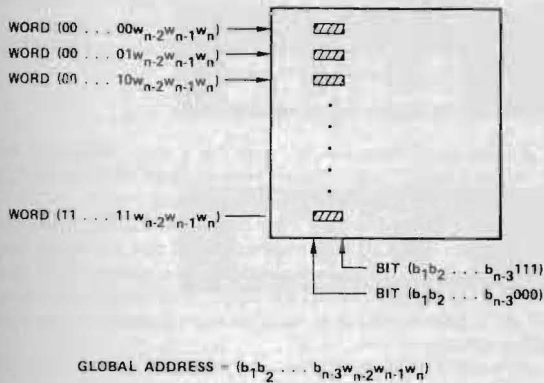


Fig. 5. Byte-access stencil.

2^{n-3} bytes in a word is accessed; the second part of the global address ($w_{n-2}w_{n-1}w_n$) specifies the subset of accessed words. Fig. 5 illustrates this byte-access stencil. Note that if a number of 2^n -byte records are stored in memory with each record occupying eight contiguous words, then the byte-access stencil accesses corresponding bytes from each record.

In general, if the access mode has m 1's and $n - m$ 0's, the access stencil covers 2^m bits of each of 2^{n-m} memory words. If the m 1's are at the right end of the access mode, the 2^m covered bits are in contiguous group, and the 2^{n-m} covered words are separated. Conversely, if the m 1's are at the

left end of the access mode, the 2^m covered bits are separated, and the 2^{n-m} covered words form a contiguous group. Where a bit of the access mode is 1, the corresponding bit of the global address is used to position the stencil "vertically" over certain words. Where an access mode bit is 0, the corresponding global address bit is used to position the access stencil "horizontally" over certain bits of words.

CONCLUSION

By storing data in a scrambled manner, doubling the address lines, and accessing memory through a scramble/unscramble network, one can use standard RAM integrated-circuits to construct a multidimensional access memory. With suitable processing elements, bit-slice access to data allows parallel vector-arithmetic and content-addressable search operations to be executed efficiently. Conventional word access also allows input, output, and scalar arithmetic to use memory efficiently.

An alternative method of achieving MDA would be skewed storage [1]; in this case, skewed storage would require one adder for each memory chip for local address generation. Scrambling the data, as shown here, greatly simplifies local address generation (n EXCLUSIVE-OR circuits versus 2^n adders).

The scramble/unscramble network can be designed to accommodate shifting and other useful permutations of data (such as those required for parallel execution of the fast Fourier transform). It can take the place of the data-routing network required by most parallel-processing architectures and therefore not add a significant cost to the total system. Thus, the cost of MDA storage need not be much higher than the cost of conventional, random-access, solid-state storage.

REFERENCES

- [1] P. Budnik and D. J. Kuck, "The organization and use of parallel memories," *IEEE Trans. Comput.*, vol. C-20, pp. 1566-1569, Dec. 1971.
- [2] T. Feng, "Data manipulating functions in parallel processors and their implementations," *IEEE Trans. Comput.*, vol. C-23, pp. 309-318, Mar. 1974.
- [3] J. D. Feldman and L. C. Fulmer, "RADCAP—An operational parallel processing facility," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 7-15.
- [4] E. W. Davis, "STARAN parallel processor system software," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 17-22.
- [5] K. E. Batchler, "STARAN parallel processor system hardware," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 405-410.
- [6] H. S. Stone, "Parallel processing with the perfect shuffle," *IEEE Trans. Comput.*, vol. C-20, pp. 153-161, Feb. 1971.




Fig. 4. MDA memory with two-access stencils.

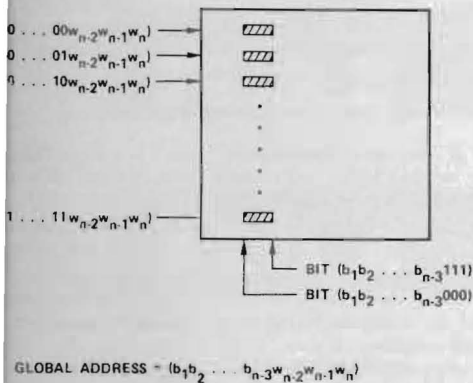


Fig. 5. Byte-access stencil.

a word is accessed; the second part of the global address specifies the subset of accessed words. Fig. 5 illustrates this stencil. Note that if a number of 2^m -byte records are memory with each record occupying eight contiguous words, the access stencil accesses corresponding bytes from each

of the access mode has m 1's and $n - m$ 0's, the access stencil of each of 2^{n-m} memory words. If the m 1's are at the right end of the access mode, the 2^m covered bits are in contiguous group, and the words are separated. Conversely, if the m 1's are at the

left end of the access mode, the 2^m covered bits are separated, and the 2^{n-m} covered words form a contiguous group. Where a bit of the access mode is 1, the corresponding bit of the global address is used to position the stencil "vertically" over certain words. Where an access mode bit is 0, the corresponding global address bit is used to position the access stencil "horizontally" over certain bits of words.

CONCLUSION

By storing data in a scrambled manner, doubling the address lines, and accessing memory through a scramble/unscramble network, one can use standard RAM integrated-circuits to construct a multidimensional access memory. With suitable processing elements, bit-slice access to data allows parallel vector-arithmetic and content-addressable search operations to be executed efficiently. Conventional word access also allows input, output, and scalar arithmetic to use memory efficiently.

An alternative method of achieving MDA would be skewed storage [1]; in this case, skewed storage would require one adder for each memory chip for local address generation. Scrambling the data, as shown here, greatly simplifies local address generation (n EXCLUSIVE-OR circuits versus 2^n adders).

The scramble/unscramble network can be designed to accommodate shifting and other useful permutations of data (such as those required for parallel execution of the fast Fourier transform). It can take the place of the data-routing network required by most parallel-processing architectures and therefore not add a significant cost to the total system. Thus, the cost of MDA storage need not be much higher than the cost of conventional, random-access, solid-state storage.

REFERENCES

- [1] P. Budnik and D. J. Kuck, "The organization and use of parallel memories," *IEEE Trans. Comput.*, vol. C-20, pp. 1566-1569, Dec. 1971.
- [2] T. Feng, "Data manipulating functions in parallel processors and their implementations," *IEEE Trans. Comput.*, vol. C-23, pp. 309-318, Mar. 1974.
- [3] J. D. Feldman and L. C. Fulmer, "RADCAP—An operational parallel processing facility," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 7-15.
- [4] E. W. Davis, "STARAN parallel processor system software," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 17-22.
- [5] K. E. Batchler, "STARAN parallel processor system hardware," in *1974 Nat. Comput. Conf., AFIPS Conf. Proc.*, vol. 43, pp. 405-410.
- [6] H. S. Stone, "Parallel processing with the perfect shuffle," *IEEE Trans. Comput.*, vol. C-20, pp. 153-161, Feb. 1971.