## COORDINATED SCIENCE LABORATORY

 APPLIED COMPUTATION THEORY GROUP
## The Cube-connected-cycles: A VERSATILE NETWORK FOR PARALLEL COMPUTATION

FRANCO P. PREPARATA JEAN VUILLEMIN

## APPROVED FOR PUBLIC RELEASE. DISTRIBUTION UNLIMITED.

| REPORT DOCUMENTATION PAGE | READ INSTRUCTIONS <br> BEFORE COMPLETING FORM |
| :--- | :--- |
| 1. REPORT NUMGER | 2. GOVT ACCESSION NO. |

## 16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)
18. SUPPLEMENTARY NOTES
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Parallel processing, VLSI design, sorting, Fourier Transform
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

We introduce an interconnection pattern of processing elements, the cube-connected-cycles (CCC), which can be used as a general purpose parallel processor. Because its design complies with present technological constraints, the CCC can also be used in the layout of many specialized large scale integrated circuits (VLSI). By combining the principles of parallelism and pipelining, the CCC can emulate the cube-connected machine and the perfect shuffle with no significant degradation of performance but with a more compact structure. We describe in detail how to program the CCC for efficiently
20. (Continued)
solving a large class of problems, which includes Fast-Fourier-Transform, sorting, permutations, and derived algorithms.

# THE CUBE-CONNECTED-CYCLES: A VERSATILE NETWORK FOR PARALLEL COMPUTATION 

## by

Franco P. Preparata and Jean Vuillemin

This work was supported in part by the National Science Foundation under Grant MCS 78-13642 and Joint Services Electronics Program under Contract N00014-79-C-0424.

> Reproduction in whole or in part is permitted for any purpose of the United States Government.

A preliminary version of this report was issued in June 1979 by I.R.I.A., Institut de Recherche d'Informatique et d'Automatique, 78150 Rocquencourt, France.

Approved for public release. Distribution unlimited.

THE CUBE-CONNECTED-CYCLES: A VERSATILE NETWORK FOR PARALLEL COMPUTATION

Franco P. Preparata Jean Vuillemin
Coordinated Science Laboratory University of Illinois Urbana, Illinois 61801

Laboratoire de Recherche en Informatique Université de Paris-Sud 91405 Orsay, France

Abstract: We introduce an interconnection pattern of processing elements, the cube-connected-cycles (CCC), which can be used as a general purpose parallel processor. Because its design complies with present technological constraints, the CCC can also be used in the layout of many specialized large scale integrated circuits (VISI). By combining the principles of parallelism and pipelining, the CCC can emulate the cube-connected machine and the perfect shuffle with no significant degradation of performance but with a more compact structure. We describe in detail how to program the CCC for efficiently solving a large class of problems, which includes Fast-Fourier-Transform, sorting, permutations, and derived algorithms.

Keywords: Parallel processing, VLSI design, sorting, Fourier Transform.

## CR Categories:

$\dagger$ This work was partially supported by National Science Foundation Grant MCS-78-13642, by the Joint Services Electronics Program Contract DAAG-29-78-C-0016, by I. R. I. A., Institut de Recherche en Informatique et Automatique, 78150 Le Chesnay, France, and by ERA 452 "al Khowarizmi" of Centre National de la Recherche Scientifique.
\# A preliminary version of this paper has been presented at the $20^{\text {th }}$ Annual Symposium on Foundations of Computer Science, Puerto Rico, Oct. 1979.

## 1. INTRODUCTION

The great technological progress embodied in very large scale integration (VLSI) of electronic circuits has made it possible to conceive large systems of processing elements cooperating in the execution of parallel algorithms. This has motivated considerable research interest in parallel computation. Unfortunately, here the situation is very different from that of serial computation, where the RAM machine [1] represents a universally accepted model. The difficulty of choosing a specific interconnection is frequently bypassed by assuming a model (shared-memory-machine) where each pair of processors is connected (or an equivalent system) [2-5]. Although not without merit, because it aims at uncovering the inherent data-dependence of given problems, this approach ignores the technological constraints of VLSI, particularly as regards the communication among the processing elements [6]. At the opposite end, other workers [7-11] suggest that processor interconnection should be limited to planar links between topologically neighboring cells (arrays or meshes). Such designs are certainly well suited for current VLSI technology, and they have cleverly been used in implementing algorithms for matrices or graph problems [9-12], for example. This type of connection, however, is not suited for efficiently implementing algorithms for various fundamental problems, such as sorting and convolution. Indeed, good algorithms for solving these problems intrinsically require data movement between processors which are topologically far apart; for example, sorting on an $n$ processor array such as ILLIAC IV requires time $\Omega(\sqrt{n})$ [8].

The purpose of the paper is to propose and analyze a new interconnection of processors, called the cube-connected-cycles, which is remarkably suited for implementing efficient algorithms such as Fast-Fourier-Transform (FFT), sorting, etc... . The geometric structure underlying the interconnections is that the k -dimensional cube. This structure which has already been studied in relation to parallel computation [13], is not readily usable for VLSI design, since each of the $2^{k}$ processors is connected to $k$ other processors.

By combining parallelism and pipelining we are able to achieve the following results:
(1) The number of connections per processor is reduced to 3.
(2) Processing time is not significantly increased with respect to that achievable on the $k$-cube structure.
(3) Programs for the individual modules are obtained in a systematic way from a standard description of the global algorithms.
(4) The overall structure complies with the basic requirements of VLSI technology: modularity, ease of layout, simplicity of communication among the processing elements, simplicity in timing and control of the entire system [14]. We also propose a wire layout of the CCC, which can be physically realized with two orthogonal layers of wires. This layout is optimal for several problems, according to a recently proposed VSLI mode1 [18].
(5) Finally we are able, without resorting to any drastic departure from classical algol-like languages, to provide fully accurate and hopefully easily understandable descriptions of our parallel programs. This is a favorable sign that parallel processing may possibly be endowed with suitable high level programming languages.

This paper is organized as follows. Section 2 introduces a class of algorithms comprising many important applications, such as merging, sorting, Fourier Transform, data rearrangement, ... . Section 3
presents models of module connections, including the CCC, allowing for efficient parallel execution of the algorithms in Section 2 . Section 4 describes the implementation of such algorithms on the CCC, and Section 5 is devoted to optimality considerations regarding a layout of the machine for VLSI realizations.

## 2. A CLASS OF HIGHLY PARALLEL ALGORITHMS

To describe our algorithms, assume that input data $t_{0}, t_{1}, \ldots, t_{n-1}$ are stored respectively in storage locations $T[0], T[1], \ldots, T[n-1]$, and that $n=2^{k}$, i.e., the number of inputs is a power of 2 . We say that an algorithm is in the DESCEND class if it performs a sequence of basic operations on data which are successively $2^{k-1}, \ldots, 2^{j}, \ldots, 2^{0}=1$ locations apart. Each basic operation $\operatorname{OPER}(\mathrm{m}, \mathrm{j} ; \mathrm{U}, \mathrm{V})$ modifies the two data items present in storage locations $U$ and $V$; the computation performed affects only the contents of $U, V$ and it may depend upon parameters $m$ and $j$, which are integers $0 \leq m<n, 0 \leq j<k$.

Algorithms in the DESCEND class are then specified as:
proc
DESCEND
for $j \leftarrow k-1$ step -1 until $j=0$
do foreach $\mathrm{m}: 0 \leq \mathrm{m}<\mathrm{n}$
pardo if bit ${ }_{j}(m)=0$ then $\operatorname{OPER}\left(m, j ; T[m], T\left[m+2^{j}\right]\right)$
fi
odpar
od
corp DESCEND.
Here, bit ${ }_{j}(m)$ is the coefficient of $2^{j}$ in the binary representation of $m=\sum_{j \geq 0}$ bit $_{j}(m) 2^{j}$. The language construct foreach $m$ : <cond (m)> pardo <action> odpar obviously indicates that all instructions <action> corresponding to values of $m$ satisfying <cond $(m)>$ can be performed simultaneously. On machines where such parallelism can be realized, DESCEND algorithms run in $k=\log _{2}(n)$ elementary steps.

We also introduce the dual class ASCEND, where the control of the algorithm is changed to

$$
\text { for } j \leftarrow 0 \text { step } 1 \text { until } j=k-1
$$

i.e., OPER is performed on data which are successively
$1=2^{0}, 2^{1}, \ldots, 2^{j}, \ldots, 2^{k-1}$ locations apart. To clarify the duality between ASCEND and DESCEND consider the binary representation of $m=\sum_{0 \leq i<k} b i t_{i}(m) \cdot 2^{i}$ and define $\tilde{m}=\sum_{0 \leq i<k} b_{i t_{i}}(m) \cdot 2^{k-i-1}$, the integer whose binary representation is the reversal of that of $m$. Once $k$ is fixed, the function: $m \rightarrow \tilde{m}$ is an involutory permutation of $0,1, \ldots, 2^{k}-1$ known as the bit reversal permutation (BRP). For example, for $k=3$, the BRP of (01234567) is (04261537).

By first applying the BRP to its inputs, an ASCEND algorithm can be transformed into a dual DESCEND algorithm (figure 1) whose basic operation $\widetilde{O P E R}$ is related to the original OPER by:

$$
\begin{aligned}
& \widetilde{\operatorname{OPER}}(\mathrm{m}, \mathrm{j} ; \mathrm{U}, \mathrm{~V})= \\
& \begin{array}{rlllllllllllllll}
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & \text { OPER }(\tilde{m}, \mathrm{k}-1-j ; \mathrm{U}, \mathrm{~V}) \\
& \text { input } & 0 & 4 & 2 & 6 & 1 & 5 & 3 & 7
\end{array}
\end{aligned}
$$

$j=2$

$$
0^{\prime} 1^{\prime} 2^{\prime} 3^{\prime} 4^{\prime} 5^{\prime} 6^{\prime} 7 \prime
$$

$j=1$
"7 "6 "5" "4" "3" "2" " " "

OPER

$$
0^{\prime} 4^{\prime} 2^{\prime} 6^{\prime} 1^{\prime} 5^{\prime} 3^{\prime} 7 \prime
$$

OPER

$$
j=0
$$

"7 "3 "5 "1" " " " " " " " "

$$
j=1
$$

$$
j=0
$$

OPER
0"'4"'2"'6"'1"'5"'3"'7"

DESCEND
Figure 1. Dual algorithms; operands are denoted by their original addresses, connecting lines show interacting operands, and priming indicates the number of operations through which an operand has been processed.

It is now time to exhibit algorithms for solving specific interesting problems. Some applications - such as bitonic merge and cyclic shift are directly within the ASCEND or DESCEND classes (simple algorithms); for these applications, all we have to do is specify $\operatorname{OPER}(m, j ; U, V)$.

Other applications (such as permutation, shuffle, unshuffle, bitreversal (BRP), odd-even-merge, Fast-Fourier-Transform, convolution, matrix transposition) have programs consisting of a short sequence of algorithms (cascaded algorithms) in the preceding class, and thus run in $0(\operatorname{logn})$ parallel steps.

We also have applications - such as bitonic sort, odd-even-sort, and calculations of symmetric functions - for which the combining step of the two results of a recursive call is itself an algorithm in one of the two preceding categories. These algorithms, which we call composite, run in $O\left((\log n)^{2}\right)$ parallel steps.

### 2.1 Bitonic Merge

The elegant algorithm for bitonic merge, due to $K$. E. Batcher [15], is ideally suited for implementation within the DESCEND class. A11 we need is to specify $\operatorname{OPER}(\mathrm{m}, \mathrm{j} ; \mathrm{U}, \mathrm{V})$ as a comparison-exchange. Precisely, in order to handle sequences which are sorted either in increasing or in decreasing order, we define ORIENTCOMPEXCHANGE ( $m, j ; \mathrm{U}, \mathrm{V}$ ) as

$$
\begin{array}{r}
\text { if bit }{ }_{j}(m)=0 \text { then }(U, V) \leftarrow(\underline{\text { min }}(U, V), \underline{\max }(U, V)) \\
\text { else }(U, V) \leftarrow(\underline{\max }(U, V), \underline{\min }(U, V))
\end{array}
$$

fi.
Batcher's odd-even merge $[15,16]$ can also be programmed as a cascaded algorithm, running in $0(\log n)$ parallel steps.

### 2.2 Radix-2 Fast-Fourier-Transforms and Convolution

The important FFT algorithm can be set in the ASCEND class. Let $\omega$ be a primitive root of unity of order $n=2^{k}$. If $<A_{0}, \ldots, A_{n-1}>$ is the Fourier Transform of vector $\left\langle a_{0}, \ldots, a_{n-1}\right\rangle$, it is well-known that $A_{j}=U_{j}+\omega^{j} V_{j}$ and $A_{j+2} k-1=U_{j}-\omega^{j} V_{j}$ where the $U^{\prime} s$ and $V^{\prime} s$
are respectively the Fourier Transforms, with primitive root $\omega^{2}$, of the "even" subsequence $<a_{0}, a_{2}, \ldots, a_{2}{ }_{-2}>$ and the "odd" subsequence $<a_{1}, a_{3}, \ldots, a_{2_{-1}}>$; we call the $\omega^{j_{1}}{ }^{2}$ the combining root powers.

The above relationships indicate that the sequence $<a_{0}, \ldots, a_{n-1}>$ must be initially rearranged by means of the bit-reversal permutation. Once the desired reconfiguration has been achieved, we may proceed with the actual FFT computation, which is in the ASCEND class.

Its basic operation $\operatorname{OPER}(\mathrm{m}, \mathrm{j} ; \mathrm{U}, \mathrm{V})$ is specified by $(U, V) \leftarrow(U+\alpha V, U-\alpha V)$ where $\alpha=\omega^{m \cdot 2^{k-j}}$.

It is not hard to show that $\alpha$ can be computed efficiently at each step; precisely, the time used by each module to compute, by successive squaring, the required combining root powers for the entire algorithm is $j\left((\log \log n)^{2}\right)=o(\log n)$. Using a sequence of two inverse Fourier transforms in the classical manner [1] allows one to compute the convolution of two sequences, from which a wealth of applications can be derived (see [1]).

### 2.3 Data Rearrangements

Being able to efficiently permute the data is obviously important for may applications. For example, the BRP rearrangement is a necessary preliminary step to the FFT algorithm of the preceding section. Some permutations, such as cyclic shifts, shuffle, and unshuffle can be computed by algorithms in ASCEND or DESCEND, as the reader will enjoy discovering for himself (here "shuffle" of ( $0,1,2, \ldots, 2^{k}-1$ ) is $\left(0,2^{k-1}, 1,2^{k-1}+1, \ldots, 2^{k-1}-1,2^{k}-1\right)$ and "unshuffle" is the inverse permutation). Other permutations, such as BRP or matrix transpose, are computed by cascaded algorithms. In general, we can emulate a Benes permutation network [21] by a sequence ASCEND;DESCEND, thus in time $O$ (logn) ; it must be pointed out, however, that to realize an arbitrary permutation, the exchange information must be precomputed.

### 2.4 Sorting and Calculation of Symmetric Functions

The previously described merge routines can be used as the basis of efficient sorting algorithms. A sequence of input keys is divided into two halves, each of which is recursively sorted (in opposite order in the case of bitonic sort), and then merged using either of the above merge routines. Both algorithms run in time $O\left((\operatorname{logn})^{2}\right)$.

One can compute symmetric functions in a completely analogous fashion: apply recursive calls to each half of the data, and compute the convolution of the two resulting sequences, again in time $O\left((\log n)^{2}\right)$.

### 2.5 Matrix Multiplications and Other Algorithms

To compute the matrix product $\mathrm{C}=\mathrm{A} \times \mathrm{B}$ of two $\mathrm{n} \times \mathrm{n}$ matrices, we must obviously first store $A=\left(A_{0}^{T} \ldots A_{n-1}^{T}\right)^{T}$ in row major order, and $B=\left(B_{0} \cdots B_{n-1}\right)$ in column major order. Assuming we have enough space and processors, i.e., $2^{k} \geq n^{3}$, we copy $A$ and $B$ into the pattern:

$$
A_{0} B_{0} A_{0} B_{1} \cdots A_{0} B_{n-1} A_{1} B_{0} \cdots A_{i} B_{j} \cdots A_{n-1} B_{n-1} \text {. Al1 this can be achieved }
$$ with simple-minded cascaded algorithms, in time $O$ (logn). Each of the scalar products $c_{i, j}=A_{i} \cdot B_{j}=\sum a_{i, k} \cdot b_{k, j}$ is computed in parallel, within $O(\operatorname{logn})$ additional time units. The results $c_{i, j}$ are then regrouped, according to the output format (say, row major).

Although the details of this algorithm are a bit tedious to describe, it should be clear that matrix multiplication can be computed in time O(logn), within our class of algorithms. In fact, a surprising number of other algorithms can be efficiently implemented within this framework, including all of the interesting algorithms for parallel processing known to the authors.
3. DESCRIPTION OF THE SCHEME

In order to efficiently implement algorithms in the DESCEND class, the most natural interconnection of modules is that of the k-dimensional
 0 to $2^{k}-1$ and is connected to each of the $k$ processors whose binary numbering differs in exactly one binary position (figure 2). Although an ASCEND or DESCEND algorithm can be implemented on such a machine in $\log _{2} n$ parallel steps, this proposal is not feasible mainly because the number $k=\log _{2} n$ of connections for each processor is too large. The unfolded $k$-cube and the perfect shuffle interconnections have been proposed [17] (figure 3), as attempts to remedy this difficulty.


Figure 2. The 3-cube.


Figure 3. Unfolded 3-cube (left) and perfect shuffle (right) interconnections.

A1though both structures have a fixed number (4) of connections per processor, their intrinsic topology make them inferior, as regards physical layout (see section 5), to the scheme we now describe.

Our parallel computing system, the cube-connected-cycles (CCC), is a network of identical processors, called modules. A module has 3 interconnection ports. Each interconnection line linking two modules can be used for the bidirectional transmission of one operand, and it is irrelevant here whether operand transmission is serial or parallel. For correctly executing the algorithms described in the preceding sections, it is indifferent to synchronize the entire system through a central clock, which defines time units for all modules, or to let synchronization
problems be settled at the level of each communication line, thus achieving a globally asynchronous system. In order to describe the interconnections, we assume for simplicity that $n$, the number of modules, is a power of two, i.e., $n=2^{k}$, and, moreover, assume that $k$ is of the form $k=r+2^{r}$; the modifications resulting when $k$ is arbitrary are straightforward (in the latter case, $r$ is the smallest integer for which $r+2^{r} \geq k$ ). Each module has a $k$-bit address $m$ which in turn is expressed as a pair ( $\ell, p$ ) of integers represented with ( $k-r$ ) and $r$ bits respectively, such that $\ell \cdot 2^{r}+p=m$.

As mentioned earlier, each module has three ports: $F$, B, and L (mnemonic for forward, backward, lateral), whose connection is entirely determined by the module address ( $\ell, \mathrm{p}$ ), that is:

$$
\begin{aligned}
& F(\ell, p) \text { is connected to } B\left(\ell,(p+1) \bmod 2^{r}\right) \\
& B(\ell, p) \text { is connected to } F\left(l,(p-1) \bmod 2^{r}\right) \\
& L(\ell, p) \text { is connected to } L\left(\ell+\varepsilon 2^{p}, p\right)
\end{aligned}
$$

where $\varepsilon=1-2$ bit $_{p}(\ell)$. The interconnection scheme is displayed in figure 4. In words, the modules are grouped into $2^{\mathrm{k}-\mathrm{r}}$ cycles, each cycle consisting of $2^{r}$ modules, cyclically connected by the $F-B$ lines. The cycles are in turn interconnected as a ( $k-r$ )-cube; if $<\mathrm{x}_{0}, \mathrm{x}_{1}, \ldots, \mathrm{x}_{\mathrm{k}-\mathrm{r}-1}>$ are the dimensions of the ( $\mathrm{k}-\mathrm{r}$ )-cube, all edges along dimension $x_{i}$, called collectively sheaf $i$, link modules whose addresses are (.,i). The total number of interconnection links is $3 \cdot 2^{\mathrm{k}-1}=\frac{3}{2} \cdot \mathrm{n}$.

Each module contains an operand register $T$, a few memory locations, and possesses basic arithmetic and logical capabilities. It is controlled by a stored program or a circuit implementation of such a program. For the time being, we make the hypothesis of unlimited parallelism,


Figure 4. The CCC interconnection scheme.
that is, the number of modules is tailored to the problem size; under this hypothesis, the one or two memories mentioned earlier suffice. Subsequently (section 4.3), under the hypothesis of limited parallelism, we shall endow each module with a small private random access memory. In either case, each module is somewhat simpler than a current microprocessor but not basically different from it.

## 4. EMULATION OF THE k -CUBE ON THE CCC

In order to implement DESCEND on the CCC, we prune the $k$-cube so as to use only connections existing in the CCC. The first stage consists in removing the sheaves corresponding to dimensions $0,1, \ldots, r-1$, and using instead the cycle connections $F$ and $B$, as introduced in section 3 . Our original DESCEND program is thus transformed to:
proc DESCEND

```
    for \(\mathrm{j} \leftarrow \mathrm{k}-1\) step -1 until \(\mathrm{j}=\mathrm{r}\)
```

    do foreach \(\mathrm{m}: 0 \leq \mathrm{m}<\mathrm{n}\)
        pardo if bit \({ }_{j}(m)=0\) then \(\operatorname{OPER}\left(m, j ; T[m], T\left[m+2^{j}\right]\right)\)
            fi
        odpar
    od;
    foreach \(\ell: 0 \leq \ell<2^{\mathrm{k}-\mathrm{r}}\) pardo LOOPOPER( \(\left.\ell\right)\) odpar
    corp DESCEND.

Here procedure DLOOPOPER( $\ell$ ) processes the data within cycle $\ell$ to compute the desired result in $O\left(2^{r}\right)$ parallel steps, as we show later. Note that the running time is still $0(k-r)+0\left(2^{r}\right)=0(\log n)$.

The second transformation consists in removing, for all $j=0, \ldots, k-r-1$, the $k-c u b e$ links pertaining to sheaf $(r+j)$, except those existing between modules whose addresses are of the form $(., j)$ : the resulting interconnection is then exactly the one of the CCC, as introduced in Section 3.

The computation corresponding to the for loop of the above algorithm can no longer be performed in one parallel step. Using repeated circular shifts within cycles, however, each operand in the
cycle can be successively brought to reside for one time unit in module $(., j)$, where $\operatorname{OPER}(., \mathrm{j} ; .,$.$) can then be executed. Although the execution$ of $\operatorname{OPER}(., j ; .,$.$) for all operands in a cycle now requires 2^{r}$ time units, this computation can be pipelined (overlapped) with the analogous operations OPER(., i;...) for $r \leq i<k$. To achieve pipelining thus requires a new function BSHIFT( $\ell$ ), which performs a cyclic backward shift of the operands in cycle $\ell$, that is:
foreach $j: 0 \leq j<2^{r}$ pardo $T\left[\ell \cdot 2^{r}+\left((j-1) \underline{\bmod } 2^{r}\right)\right] \leftarrow T\left[\ell \cdot 2^{r}+j\right]$
odpar.
The final version of DESCEND is thus:
proc DESCEND
for $i \leftarrow 2^{r}-1$ step-1 until $i=-2^{r}$
do foreach $\ell$ : $0 \leq \ell<2^{\mathrm{k}-\mathrm{r}}$
pardo foreach $\mathrm{p}: \underline{\max }(\mathrm{i}, 0) \leq \mathrm{p}<\underline{\min }\left(2^{\mathrm{r}}, 2^{\mathrm{r}}+\mathrm{i}\right)$
pardo if bit $_{\mathrm{p}}(\ell)=0$ then $\operatorname{OPER}(\mathrm{a}, \mathrm{b} ; \mathrm{U}, \mathrm{V})$
where $a=\ell \cdot 2^{r}+\left((p+i-1) \bmod 2^{r}\right)$,

$$
\mathrm{b}=\mathrm{p}+\mathrm{r},
$$

$$
\mathrm{U}=\mathrm{T}\left[\ell \cdot 2^{\mathrm{r}}+\mathrm{p}\right]
$$

$$
\mathrm{V}=\mathrm{T}\left[\left(\ell+2^{\mathrm{P}}\right) \cdot 2^{\mathrm{r}}+\mathrm{p}\right] .
$$

fi
odpar;
BSHIFT( $\ell$ ) Comment backwards shift of cycle $\ell$;
od;
Comment end of treatment on sheaves $\mathrm{k}-1, \mathrm{k}-2, \ldots, \mathrm{r}$;
foreach $\ell: 0 \leq \ell<2^{\mathrm{k}-\mathrm{r}}$ pardo LOOPOPER ( $\ell$ ) odpar
corp DESCEND.

The inner operation of the for loop is executed in two time units; one for OPER, then one for BSHIFT. The total running time is thus $4 \cdot 2^{r}$ plus the time for executing LOOPOPER. If we can ensure that LOOPOPER can be processed in time linear in the cycle size, the entire procedure will be executed on the CCC in time $O$ (logn).

Figure 5 provides a schematic view of DESCEND on the CCC, and conventions used are those of figure 1, which depicts DESCEND on the $k$-cube. Here we assume $k=3$, thus the CCC consists of 4 cycles of 1ength 2 .

## data


$1^{\prime \prime} \longleftrightarrow 0^{\prime} \quad 3^{\prime \prime} \longleftrightarrow 2^{\prime} \quad 5^{\prime \prime} \longleftrightarrow 4^{\prime} \quad 7^{\prime \prime} \longleftrightarrow 6^{\prime}$


0'" $1^{\prime \prime \prime} 2 \prime \prime$ 3'" $4 \prime \prime$ 5'" 6 '" $7 \prime \prime$


SHIFT


Figure 5. DESCEND on the CCC, $k=3$

### 4.1 Computation Within the Cycles

The next question to be addressed is the implementation of LOOPOPER ( $\ell$ ), so that it runs in time linear in the cycle length. Obviously, we are constrained to using only the $F$ and $B$ cycle links existing in the CCC. Our objective is to emulate, on the cycle of length $2^{r}$, the operation OPER as it would be executed on hypothetical r-cube sheaves. Since OPER may take place in the cycle only between adjacent modules, particular care must be exercised to ensure that the desired adjacencies, corresponding to all sheaves, be globally realized in time linear in the cycle length. The key permutations for this task are based on the perfect unshuffle $[16,17]$. Specifically, UNSHUFFLE ( $\ell, i$ ) performs the perfect-unshuffle operation on each of the $2^{\mathrm{r}-\mathrm{i}-1}$ contiguous blocks of length $2^{i+1}$ into which $\mathrm{T}\left[\ell \cdot 2^{\mathrm{r}}::(\ell+1) \cdot 2^{\mathrm{r}}-1\right]$ is subdivided, and is realized as follows:
proc UNSHUFFLE ( $\ell, \mathrm{i}$ )

$$
\text { for } b \leftarrow 2^{i} \text { step }-1 \text { until } b=2
$$

$$
\text { do foreach } \mathrm{m}: \mathrm{m}=\ell \cdot 2^{\mathrm{r}}+(2 \cdot \mathrm{~s}+1) \cdot 2^{i}+\mathrm{p}
$$

$$
\text { where } 0 \leq \mathrm{s}<2^{\mathrm{r}-\mathrm{i}-1},-\mathrm{b}<\mathrm{p}<\mathrm{b} \text {, }
$$

$$
(p \underline{\bmod } 2)=(b \underline{\bmod } 2)
$$

$$
\text { pardo } \mathrm{T}[\mathrm{~m}-1] \leftrightarrow \mathrm{T}[\mathrm{~m}] \text { odpar }
$$

## od

corp UNSHUFFLE.
Clearly, UNSHUFFLE $(\ell, i)$ runs in $\left(2^{i}-1\right)$ parallel step. It is also easy to realize that the program
proc BRP(l)
for $i \leftarrow r-1$ step -1 until $i=1$ do UNSHUFFLE $(\ell, i)$ od corp BRP
realizes the bit-reversal permutation of $T\left[\ell \cdot 2^{r}::(\ell+1) 2^{r}-1\right]$ with reference to the $r$ least-significant bits of the addresses.

We can now elucidate the general format of LOOPOPER, which consists of a sequence of unshuffle-operation pairs, each emulating a sheaf operation. This is preceded by BRP, so that, upon completion, the results are in the correct order (see figure 6). In the description below the parameter a gives the original address of the operand which is brought to module ( $\ell, \mathrm{p}$ ) by the sequence: BRP; UNSHUFFLE $(\ell, 0)$;
$\operatorname{UNSHUFFLE}(\ell, 1) ; \ldots ; \operatorname{UNSHUFFLE}(\ell, \mathrm{r}-1-\mathrm{j})$. (Recall that $\tilde{\mathrm{q}}$ denotes the integer whose binary representation is the reversal of that of the integer q.)
proc LOOPOPER( $\ell$ )
$\operatorname{BRP}(\ell)$;
for $\mathrm{j} \leftarrow \mathrm{r}-1$ step -1 until $\mathrm{j}=0$
do foreach $\mathrm{q}: 0 \leq \mathrm{q}<2^{\mathrm{r}}$, bit ${ }_{0}(\mathrm{q})=0$
pardo $\operatorname{OPER}\left(\mathrm{a}, \mathrm{j} ; \mathrm{T}\left[\ell \cdot 2^{\mathrm{r}}+\mathrm{q}\right], \mathrm{T}\left[\ell \cdot 2^{\mathrm{r}}+\mathrm{q}+1\right]\right)$ where $a=\ell \cdot 2^{r}+\left(\tilde{q} \bmod 2^{j}\right)+\left(q \bmod 2^{r-j}\right) \cdot 2^{j}$. odpar;
UNSHUFFLE ( $\ell, \mathrm{j}$ )
od
corp LOOPOPER.


Figure 6. A schematic presentation of LOOPOPER for $r=3$.

With respect to execution time, we noted that $\operatorname{UNSHUFFLE}(\cdot, i)$ runs in time $O\left(2^{i}\right)$; thus BRP and LOOPOPER jointly run in $0\left(1+2+2^{2}+\ldots+2^{r-1}\right)=0\left(2^{r}\right)$ steps, linear in the cycle length.

### 4.2 Programs for each Module of the CCC

From the preceding global description of DESCEND, it is rather straightforward to produce the sequential program of module ( $\ell, \mathrm{p}$ ). The program MODULE ( $\ell, \mathrm{p}$ ) for a given DESCEND algorithm is of the form: HIGHSHEAVES ( $\ell, \mathrm{p}$ ) ; LOWSHEAVES ( $\ell, \mathrm{p}$ ), which respectively implement the ( $k-r$ )-cube operation and LOOPOPER. The entire MODULE ( $l, p$ ) is of a very simple nature: it basically counts up time and, at each time unit numbered $t$, it tests a simple logical condition involving $\ell, p$, and $t$; depending on this test, either it does nothing, or it exchanges operands, or it exchanges operands and performs an operation on them. The details of these programs are omitted for the sake of brevity.

The precise execution time of DESCEND (or ASCEND) on the CCC is given by the formula:

$$
T=4.2^{r} \cdot T_{C C C}+\left(r+2^{r}\right) T_{\text {oper }}
$$

where $T_{C C C}$ is the time required for stepping up the control variable $t$, testing it and performing one data exchange on some of the links; $T_{\text {oper }}$ is the time required for computing $\operatorname{OPER}(\mathrm{m}, \mathrm{j} ; \mathrm{U}, \mathrm{V})$ within each module.

### 4.3 Limited Parallelism

So far, we have assumed that the size $n$ of the CCC was tailored to the application. To cope with the realistic situation where the number $N$ of inputs is larger than the size $n$ of the CCC, we suggest to let each module of the CCC be a full fledged microprocessor endowed with a private RAM memory.

Assuming for simplicity that $N=s n$, with $s=2^{q}$ integer, we require that the RAM memory of each module be of size $s$ and denote by $T\left[m, 0:: 2^{q}-1\right]$ the private memory locations of module $m$. The input $a_{0}, \ldots, a_{N-1}$ is divided into consecutive blocks of size $s$, each block being stored within a module of the CCC, so that $T[m, j]=a$ $2^{q} \cdot \mathrm{~m}+\mathrm{j}$ for $0 \leq j<2^{q}$.

The only modification concerns the program MODULE ( $\ell, \mathrm{p}$ ) (see Section 4.2), which now assumes the format HIGHSHEAVES ( $\ell, \mathrm{p}$ ); LOWSHEAVES ( $\ell, \mathrm{p}$ ); LOCAL $(\ell, p)$. Programs for HIGHSHEAVES and LOWSHEAVES are the same as before, except that each operation and data transmission is now successively performed on the $2^{q}$ data items of each module. As for LOCAL:
proc $\operatorname{LOCAL}(\ell, p)$
$u \leftarrow m \cdot 2^{q}$
for $j \leftarrow q-1$ step -1 until $j=0$
do for $i-0$ step 1 until $i=2^{q}-1$
do if $^{\text {bit }}{ }_{j}(i)=0$
then $\operatorname{OPER}\left(u+i, j ; T[m, i], T\left[m, i+2^{j}\right]\right)$ fi
od
od
corp LOCAL.
It should be clear by now that all of the algorithms described in Section 1 can be applied here. A direct analysis shows that, on a CCC consisting of $n$ processors, each processor having memory $\frac{N}{n}$, we can process $N$ inputs in time $O\left(\frac{N}{n} \cdot \log N\right)$ for algorithms in the classes ASCEND or

DESCEND, thus achieving the optimal speed-up possible with $n$ processors.

## 5. LAYOUT OF THE CCC FOR VLSI

It is interesting to examine the just described CCC within the framework of the "VLSI model of computation" recently proposed [14, 18, 19]. In this model, each wire has unit width on the silicon chip and transmits a unit of information in a unit of time; information is taken from, or delivered to, special areas on the chip, called nexuses, each associated with a module. Within this model, which takes realistic account of the placement of modules and interconnection, C. D. Thompson has studied the implementation of the Fast-Fourier-Transform [18] and has elucidated significant relationships between input size $n$, chip area $A$, processing time $T$, and the so-called minimal bisection width $\omega$. ${ }^{(1)}$ Thompson has shown that $A \geq \omega^{2} / 4$ in general, and that, for the $n$-point FFT, $T \geq n / 2 \omega$, thus establishing the bound $\mathrm{AT}^{2} \geq \mathrm{n}^{2} / 16$. The lower bound for time applies to a wider class of problems, as shown by the following proposition which we state without proof:

Proposition: In the VLSI model (Thompson [18]), time $T \geq \frac{\mathrm{n}}{2 \omega}$ is required to merge two sorted sequences of length $n / 2$, or to realize the data rearrangement specified by some permutation drawn from a transitive group of permutations.

As a consequence, we have $A T^{2} \geq \frac{\mathrm{n}^{2}}{16}$ for all such problems.
${ }^{(1)}$ For a graph $G=(V, E)$ the minimal bisection width $\omega$ is defined as the smallest integer such that $w=\left|\left\{(u, v) \in E: u \in v_{1}, v \in v_{2}\right\}\right|$, where $\left\{\mathrm{v}_{1}, \mathrm{v}_{2}\right\}$ is a partition of v with $\left|\mathrm{v}_{1}\right| \leq\left|\mathrm{v}_{2}\right| \leq\left|\mathrm{v}_{1}\right|+1$.
${ }^{(2)} A$ subgroup $G$ of the symmetric group $S_{n}$ is said to be transitive if $\forall i, j 1 \leq i, j \leq n$, $\operatorname{Ha}_{\sigma} \in G: \sigma(i)=j$, meaning that data located in any position of the machine may be moved into any other position of the machine.

With the CCC, we have shown that operations such as FFT, merging, cyclic shifts, shuffles, etc., are all realizable in the minimal achievable time $T=O(\operatorname{logn})$. We now demonstrate that $A=O\left(n^{2} / \operatorname{logn}{ }^{2}\right)$ thus achieving the lower bound exactly; this means that the CCC is optimal in the VLSI mode1 for FFT , merging of sorted sequences, and realization of permutations drawn from a transitive group. In contrast, known layouts for the $k$-cube or the perfect shuffle have area of a larger order.

To achieve $A=O\left((n / \log n)^{2}\right)$ for the CCC, consider a layout which uses two sheaves of evenly spaced wires, horizontal and vertical, used respectively for cube and cycle connections. Figure 7 pictorially provides base, inductive hypothesis, and extension, to prove that an $n=s \cdot 2^{s}$ module CCC can be placed on a $2^{s} \times\left(2.2^{s}-1\right)$ chip; since $s \simeq \log _{2}\left(n / \log _{2} n\right)$, the chip size is about $\left(n / \log _{2} n\right) \times\left(2 n / \log _{2} n-1\right)=0\left((n / \log n)^{2}\right)$. Slightly more complicated constructions yield somewhat more efficient module placements as suggested by figure 8.


Figure 7. A standard layout for the interconnection of $4.2^{4}$ modules.


Figure 8. A more economical layout for the interconnection of $4.2^{4}$ modules.

For pedagogical reasons, the CCC introduced so far has a number $n=s \cdot 2^{s}$ of processing modules with $s=2^{r}$ a power of 2 . A more general version of the CCC can be designed, comprising $n=h \cdot 2^{s}$ modules. Each of the $2^{s}$ cycles of the machine has $h \geq s$ modules. The lower $s \times 2^{s}$ modules of the cycles exhibit the horizontal interconnection of standard CCC, while the ( $\mathrm{h}-\mathrm{s}$ ) $\times 2^{\mathrm{s}}$ higher modules only have vertical (cycle) connections, as indicated in figure 9. Such a layout has height $2^{s}+h-s$ and width $2^{s+1}$ (in unit wire width). The programs presented in section 4 can be adapted to run on such a machine by simply ignoring operations pertaining to non-existing horizontal (external) links, and their running time is proportional to the cycle length $h$. We see that, for any value of $h$ satisfying $\log _{2} n \leq h \leq \sqrt{n}$, the area $\times(\text { time })^{2}$ product

$$
A T^{2}=\left(\frac{n}{h}+h-\log \left(\frac{n}{h}\right)\right) \times \frac{n}{h} \times h^{2}=n^{2}+n h^{2}-n h \log \left(\frac{n}{h}\right)=0\left(n^{2}\right)
$$

meets the optimal theoretical bound, to within a constant factor. Of particular interest is the choice $h=0(\sqrt{n})$, which leads to a running time $T=0(\sqrt{n})$ and uses the minimal achievable area $A=0(n)$.


Figure 9. A standard layout for an $h \times 2^{s} \operatorname{CCC}(h=6, s=4)$.

## 6. CONCLUSION

In this paper, we have proposed a structure which can be used for direct hardware implementation of specific useful algorithms, or, as suggested in section 4.3 , as a general purpose parallel processing system.

We expect the CCC to be practically feasible in the present state of the technology, and to be capable of executing efficiently a wide variety of algorithms. The extent of the class of algorithms amenable to efficient CCC processing is not yet well understood, but it goes beyond the applications described in Section 1 ; in particular, it includes a variety of matrix and graph algorithms, as well as arithmetic and algebraic problems.

Another salient feature of this work is the possibility which appears to exist of developing a high level, general purpose language for parallel programming, which would nevertheless be automatically compilable on systems such as the CCC.

## REFERENCES

[1] A. V. Aho, J. E. Hopcroft and J. D. Ullman, The Analysis and Design of Computer Algorithms. Addison-Wesley, Reading, Mass., 1974.
[2] D. Heller, "A survey of parallel algorithms in numerical linear algebra," Dept. of Comp. Sci., Carnegie-Mellon University, Pittsburgh, Pa., Feb. 1976.
[3] D. S. Hirschberg, "Fast parallel sorting algorithms," Communications of the ACM, vol. 21, no. 8, pp. 657-661, 1978.
[4] L. G. Valiant, "Parallelism in comparison problems," SIAM Journal on Computing, vol. 4, 3, pp. 348-355, Sept. 1975.
[5] F. P. Preparata, "New parallel sorting schemes," IEEE Transactions on Computers, vol. C-27, no. 7, pp. 669-673, July 1978.
[6] W. M. Gentleman, "Some complexity results for matrix computation on paralle1 processors," Journal of the ACM, 25, 1, pp. 112-115, Jan. 1978.
[7] G. H. Barnes et al, "The ILLIAC IV computer," IEEE Transactions on Computers, vol. C-17, pp. 746-757, 1968.
[8] C. D. Thompson and H. T. Kung, "Sorting on a mesh connected computer," Proc. ACM-SIGACT Symp. on Theory of Computing, Hershey, Pa., pp. 58-64, May 1976.
[9] H. T. Kung and C.E. Leiserson, "Algorithms for VLSI processor arrays," Symposium on Sparse Matrix Computations, Knoxville, Tenn., Nov. 1978.
[10] D. Nassimi and S. Sahni, "Bitonic sort on a mesh-connected parallel computer," IEEE Transactions on Computers, vo1. C-28, no. 1, pp. 2-7, Jan. 1979.
[11] L. J. Guibas, H. T. Kung and C. D. Thompson, "Direct VLSI implementation of combinatorial algorithms," Research Report, Dept. of Comp. Sci., Carnegie-Mellon University, Pittsburgh, Pa., March 1979.
[12] K. N. Levitt and W. H. Kautz, "Cellular arrays for the solution of graph problems," Comm. of the ACM, vol. 15, no. 9, pp. 789-801, 1972.
[13] M. C. Pease, "The indirect binary n-cube microprocessor array," IEEE Transactions on Computers, vo1. C-26, no. 5, pp. 458-473, May 1977.
[14] A. M. Mead and L. A. Convay, Introduction to VLSI Systems. Textbook in preparation (1979).
[15] K. E. Batcher, "Sorting networks and their applications," Proc. AFIPS Spring Joint Computer Conference, vol. 32, pp. 307-314, April 1968.
[16] D. E. Knuth, The Art of Computer Programming Volume 3, Sorting and Searching. Addison-Wesley, Reading, Mass., 1973.
[17] H. S. Stone, "Parallel processing with the perfect shuffle," IEEE Transactions on Computers, vo1. C-20, pp. 153-161, 1971.
[18] C. D. Thompson, "Area-time complexity for VLSI," Proc. of the $11^{\text {th }}$ Annual ACM Symp. of Theory of Computing, pp. 81-88, 1979.
[19] C. D. Thompson, "A complexity theory for VLSI," Ph.D. Thesis, Carnegie-Mellon University, Dept. of Comp. Sci., 1980.
[20] M. C. Pease, "An adaptation of the Fast Fourier transform for paralle1 processing," Journal of the ACM, $15(2)$, pp. 252-264, April 1968.
[21] A Waksman, "A permutation network," Journal of the ACM, 15 (1), pp. 159-163 (1968).

