

Chapter 8

Parallel Recursion

8.1 Parallelism and Recursion

Many important synchronous parallel algorithms—Fast Fourier Transform, routing and permutation, Batcher sorting schemes, solving tridiagonal linear systems by odd-even reduction, prefix-sum algorithms—are conveniently formulated in a recursive fashion. The network structures on which parallel algorithms are typically implemented—butterfly, sorting networks, hypercube, complete binary tree—are, also, recursive in nature. However, parallel recursive algorithms are typically described iteratively, one parallel step at a time¹. Similarly, the connection structures are often explained pictorially, by displaying the connections between one “level” and the next. The mathematical properties of the algorithms and connection structures are rarely evident from these descriptions.

A data structure, *powerlist*, is proposed in this paper that highlights the role of both parallelism and recursion. Many of the known parallel algorithms—FFT, Batcher Merge, prefix sum, embedding arrays in hypercubes, etc.—have surprisingly concise descriptions using powerlists. Simple algebraic properties of powerlists permit us to deduce properties of these algorithms employing structural induction.

8.2 Powerlist

The basic data structure on which recursion is employed (in LISP[37] or ML[38]) is a *list*. A list is either empty or it is constructed by concatenating an element to a list. (We restrict ourselves to finite lists throughout this paper.) We call such a list *linear* (because the list length grows by 1 as a result of applying the basic constructor). Such a list structure seems unsuitable for expressing parallel algorithms succinctly; an algorithm that processes the list elements has to describe how successive elements of the list are processed.

¹A notable exception is the recursive description of a prefix sum algorithm in [26].

We propose *powerlist* as a data structure that is more suitable for describing parallel algorithms. The base—corresponding to the empty list for the linear case—is a list of one element. A longer powerlist is constructed from the elements of two powerlists of the same length, as described below. Thus, a powerlist is multiplicative in nature; its length doubles by applying the basic constructor.

There are two different ways in which powerlists are joined to create a longer powerlist. If p, q are powerlists of the same length then

$p \mid q$ is the powerlist formed by concatenating p and q , and
 $p \bowtie q$ is the powerlist formed by successively taking alternate items from p and q , starting with p .

Further, we restrict p, q to contain similar elements (defined in Section 8.2.1).

In the following examples the sequence of elements of a powerlist are enclosed within angular brackets.

$$\begin{aligned} \langle 0 \rangle \mid \langle 1 \rangle &= \langle 0 \ 1 \rangle \\ \langle 0 \rangle \bowtie \langle 1 \rangle &= \langle 0 \ 1 \rangle \\ \langle 0 \ 1 \rangle \mid \langle 2 \ 3 \rangle &= \langle 0 \ 1 \ 2 \ 3 \rangle \\ \langle 0 \ 1 \rangle \bowtie \langle 2 \ 3 \rangle &= \langle 0 \ 2 \ 1 \ 3 \rangle \end{aligned}$$

The operation \mid is called *tie* and \bowtie is *zip*.

8.2.1 Definitions

A data item from the linear list theory will be called a *scalar*. (Typical scalars are the items of base types—integer, boolean, etc.—tuples of scalars, functions from scalars to scalars and linear lists of scalars.) Scalars are uninterpreted in our theory. We merely assume that scalars can be checked for type compatibility. We will use several standard operations on scalars for purposes of illustration.

Notational Convention : Linear lists will be enclosed within square brackets, $[]$.

A *powerlist* is a list of length 2^n , for some n , $n \geq 0$, all of whose elements are similar. We enclose powerlists within angular brackets, $\langle \rangle$.

Two scalars are *similar* if they are of the same type. Two powerlists are *similar* if they have the same length and any element of one is similar to any element of the other. (Observe that *similar* is an equivalence relation.)

Let S denote an arbitrary scalar, P a powerlist and u, v similar powerlists. A recursive definition of a powerlist is

$$\langle S \rangle \text{ or } \langle P \rangle \text{ or } u \mid v \text{ or } u \bowtie v$$

Examples




Figure 8.1: Representation of a complete binary tree where the data are at the leaves. For leaf nodes, the powerlist has one element. For non-leaf nodes, the powerlist has two elements, namely, the powerlists for the left and right subtrees.

- <2> powerlist of length 1 containing a scalar
- <<2>> powerlist of length 1 containing a powerlist of length 1 of scalar
- <> not a powerlist
- <[]> powerlist of length 1 containing the empty linear list
- <<[2] [3 4 7]> <<[4] []>>
powerlist of length 2, each element of which is a powerlist of length 2, whose elements are linear lists of numbers
- <<0 4> <1 5> <2 6> <3 7>>
a representation of the matrix $\begin{bmatrix} 0 & 1 & 2 & 3 \\ 4 & 5 & 6 & 7 \end{bmatrix}$ where each column is an element of the outer powerlist.
- <<0 1 2 3> <4 5 6 7>>
another representation of the above matrix where each row is an element of the outer powerlist.
- <<<a> > <<c> <d>>>
a representation of the tree in Figure 8.1. The powerlist contains two elements, one each for the left and right subtrees.

8.2.2 Functions over Powerlists

Convention : We write function application without parentheses where no confusion is possible. Thus, we write “ $f x$ ” instead of “ $f(x)$ ” and “ $g x y$ ” instead of “ $g(x, y)$ ”. The constructors $|$ and \bowtie have the same binding power and their binding power is lower than that of function application. Throughout this paper, S denotes a scalar, P a powerlist and x, y either scalar or powerlist. Typical names for powerlist variables are p, q, r, s, t, u, v . \square

Functions over linear lists are typically defined by case analysis—a function is defined over the empty list and, recursively, over non-empty lists. Functions over powerlists are defined analogously. For instance, the following function, rev , reverses the order of the elements of the argument powerlist.

$$\begin{aligned} rev(x) &= \langle x \rangle \\ rev(p | q) &= (rev q) | (rev p) \end{aligned}$$

The case analysis, as for linear lists, is based on the length of the argument powerlist. We adopt the pattern matching scheme of ML[38] and Miranda[52]² to *deconstruct* the argument list into its components, p and q , in the recursive case. Deconstruction, in general, uses the operators $|$ and \bowtie ; see Section 8.3. In the definition of rev , we have used $|$ for deconstruction; we could have used \bowtie instead and defined rev in the recursive case by

$$\text{rev}(p \bowtie q) = (\text{rev } q) \bowtie (\text{rev } p)$$

It can be shown, using the laws in Section 8.3, that the two proposed definitions of rev are equivalent and that

$$\text{rev}(\text{rev } P) = P$$

for any powerlist P .

Scalar Functions

Operations on scalars are outside our theory. Some of the examples in this paper, however, use scalar functions, particularly, addition and multiplication (over complex numbers) and *cons* over linear lists. A scalar function, f , has zero or more scalars as arguments and its value is a scalar. We coerce the application of f to a powerlist by applying f “pointwise” to the elements of the powerlist. For a scalar function f of one argument we define

$$\begin{aligned} f\langle x \rangle &= \langle f \ x \rangle \\ f(p \mid q) &= (f \ p) \mid (f \ q) \end{aligned}$$

It can be shown that

$$f(p \bowtie q) = (f \ p) \bowtie (f \ q)$$

A scalar function that operates on two arguments will often be written as an infix operator. For any such function \oplus and similar powerlists p, q, u, v , we have

$$\begin{aligned} \langle x \rangle \oplus \langle y \rangle &= \langle x \oplus y \rangle \\ (p \mid q) \oplus (u \mid v) &= (p \oplus u) \mid (q \oplus v) \\ (p \bowtie q) \oplus (u \bowtie v) &= (p \oplus u) \bowtie (q \oplus v) \end{aligned}$$

Thus, scalar functions commute with both $|$ and \bowtie .

Note : Since a scalar function is applied recursively to each element of a powerlist, its effect propagates through all “levels”. Thus, $+$ applied to matrices forms their elementwise sum. \square

²Miranda is a trademark of Research Software Ltd.

8.2.3 Discussion

The base case of a powerlist is a singleton list, not an empty list. Empty lists (or, equivalent data structures) do not arise in the applications we have considered. For instance, in matrix algorithms the base case is a 1×1 matrix rather than an empty matrix, Fourier transform is defined for a singleton list (not the empty list) and the smallest hypercube has one node.

The recursive definition of a powerlist says that a powerlist is either of the form $u \bowtie v$ or $u \mid v$. In fact, every non-singleton powerlist can be written in either form in a unique manner (see Laws in Section 8.3). A simple way to view $p \mid q = L$ is that if the elements of L are indexed by n -bit strings in increasing numerical order (where the length of L is 2^n) then p is the sublist of elements whose highest bit of the index is 0 and q is the sublist with 1 in the highest bit of the index. Similarly, if $u \bowtie v = L$ then u is the sublist of elements whose lowest bit of the index is 0 and v 's elements have 1 as the lowest bit of the index.

At first, it may seem strange to allow two different ways for constructing the same list—using tie or zip. As we see in this paper this causes no difficulty, and further, this flexibility is essential because many parallel algorithms—the Fast Fourier Transform being the most prominent—exploit both forms of construction.

We have restricted u, v in $u \mid v$ and $u \bowtie v$ to be similar. This restriction allows us to process a powerlist by recursive divide and conquer, where each division yields two halves that can be processed in parallel, by employing the same algorithm. (Square matrices, for instance, are often processed by quartering them. We will show how quartering, or quadrupling, can be expressed in our theory.) The similarity restriction allows us to define complete binary trees, hypercubes and square matrices that are not “free” structures.

The length of a powerlist is a power of 2. This restricts our theory somewhat. It is possible to design a more general theory eliminating this constraint; we sketch an outline in Section 8.6.

8.3 Laws

L0. For singleton powerlists, $\langle x \rangle, \langle y \rangle$

$$\langle x \rangle \mid \langle y \rangle = \langle x \rangle \bowtie \langle y \rangle$$

L1. (Dual Deconstruction)

For any non-singleton powerlist, P , there exist similar powerlists r, s, u, v such that

$$P = r \mid s \text{ and } P = u \bowtie v$$

L2. (Unique Deconstruction)

$$\begin{aligned} \langle x \rangle = \langle y \rangle &\equiv (x = y) \\ (p \mid q = u \mid v) &\equiv (p = u \wedge q = v) \\ (p \bowtie q = u \bowtie v) &\equiv (p = u \wedge q = v) \end{aligned}$$

$$\begin{aligned} \text{L3. (Commutativity of } | \text{ and } \bowtie) \\ (p | q) \bowtie (u | v) = (p \bowtie u) | (q \bowtie v) \end{aligned}$$

These laws can be derived by suitably defining tie and zip, using the standard functions from the linear list theory. One possible strategy is to define tie as the concatenation of two equal length lists and then, use the Laws L0 and L3 as the definition of zip; Laws L1, L2 can be derived next. Alternatively, these laws may be regarded as axioms relating tie and zip.

Law L0 is often used in proving base cases of algebraic identities. Laws L1, L2 allow us to uniquely deconstruct a non-singleton powerlist using either $|$ or \bowtie . Law L3 is crucial. It is the only law relating the two construction operators, $|$ and \bowtie , in the general case. Hence, it is invariably applied in proofs by structural induction where both constructors play a role.

Inductive Proofs

Most proofs on powerlists are by induction on the length, depth or shape of the list. The length, len , of a powerlist is the number of elements in it. Since the length of a powerlist is a power of 2, the logarithmic length, lgl , is a more useful measure. Formally,

$$\begin{aligned} \text{lgl}\langle x \rangle &= 0 \\ \text{lgl}(u | v) &= 1 + (\text{lgl } u) \end{aligned}$$

The *depth* of a powerlist is the number of “levels” in it.

$$\begin{aligned} \text{depth } \langle S \rangle &= 0 \\ \text{depth } \langle P \rangle &= 1 + (\text{depth } P) \\ \text{depth } (u | v) &= \text{depth } u \end{aligned}$$

(In the last case, since u, v are similar powerlists they have the same depth.) Most inductive proofs on powerlists order them lexicographically on the pair (depth, logarithmic length). For instance, to prove that a property Π holds for all powerlists, it is sufficient to prove

$$\begin{aligned} \Pi\langle S \rangle, \text{ and} \\ \Pi P \Rightarrow \Pi\langle P \rangle, \text{ and} \\ (\Pi u) \wedge (\Pi v) \wedge (u, v) \text{ similar} \Rightarrow \Pi(u | v) \end{aligned}$$

The last proof step could be replaced by

$$(\Pi u) \wedge (\Pi v) \wedge (u, v) \text{ similar} \Rightarrow \Pi(u \bowtie v)$$

The *shape* of a powerlist P is a sequence of natural numbers n_0, n_1, \dots, n_d where d is the depth of P and

$$\begin{aligned} n_0 &\text{ is the logarithmic length of } P, \\ n_1 &\text{ is the logarithmic length of (any) element of } P, \text{ say } r \\ n_2 &\text{ is the logarithmic length of any element of } r, \dots \\ &\vdots \end{aligned}$$

A formal definition of shape is similar to that of depth. The shape is a linear sequence because all elements, at any level, are similar. The shape and the type of the scalar elements define the structure of a powerlist completely. For inductive proofs, the powerlists may be ordered lexicographically by the pair (depth, shape), where the shapes are compared lexicographically.

Example : The *len*, *lgl* and *depth* of $\langle \langle 0\ 1\ 2\ 3 \rangle \langle 4\ 5\ 6\ 7 \rangle \rangle$ are, 2, 1, 1, respectively. The *shape* of this powerlist is the sequence, 1 2, because there are 2 elements at the outer level and 4 elements at the inner level.

8.4 Examples

We show a few small algorithms on powerlists. These include such well-known examples as the Fast Fourier Transform and Batchier sorting schemes. We restrict the discussion in this section to simple (unnested) powerlists (where the depth is 0); higher dimensional lists (and algorithms for matrices and hypercubes) are taken up in a later section. Since the powerlists are unnested, induction based on length is sufficient to prove properties of these algorithms.

8.4.1 Permutations

We define a few functions that permute the elements of powerlists. The function *rev*, defined in Section 8.2.2, is a permutation function. These functions appear as components of many parallel algorithms.

Rotate

Function *rr* rotates a powerlist to the right by one; thus, $rr\langle a\ b\ c\ d \rangle = \langle d\ a\ b\ c \rangle$. Function *rl* rotates to the left: $rl\langle a\ b\ c\ d \rangle = \langle b\ c\ d\ a \rangle$.

$$\begin{aligned} rr\langle x \rangle &= \langle x \rangle & , & & rl\langle x \rangle &= \langle x \rangle \\ rr(u \bowtie v) &= (rr\ v) \bowtie u & , & & rl(u \bowtie v) &= v \bowtie (rl\ u) \end{aligned}$$

There does not seem to be any simple definition of *rr* or *rl* using $|$ as the deconstruction operator. It is easy to show, using structural induction, that *rr*, *rl* are inverses. An amusing identity is $rev(rr(rev(rr\ P))) = P$.

A powerlist may be rotated through an arbitrary amount, k , by applying k successive rotations. A better scheme for rotating $(u \bowtie v)$ by k is to rotate both u , v by about $k/2$. More precisely, the function *grr* (given below) rotates a powerlist to the right by k , where $k \geq 0$. It is straightforward to show that for all $k, k \geq 0$, and all p , $(grr\ k\ p) = (rr^{(k)}\ p)$, where $rr^{(k)}$ is the k -fold application of *rr*.

$$\begin{aligned} grr\ k\ \langle x \rangle &= \langle x \rangle \\ grr\ (2 \times k)\ (u \bowtie v) &= (grr\ k\ u) \bowtie (grr\ k\ v) \\ grr\ (2 \times k + 1)\ (u \bowtie v) &= (grr\ (k + 1)\ v) \bowtie (grr\ k\ u) \end{aligned}$$

P 's indices	=	(000	001	010	011	100	101	110	111)
List P	=	$\langle a$	b	c	d	e	f	g	$h \rangle$
P 's indices rotated right	=	(000	100	001	101	010	110	011	111)
$rs P$	=	$\langle a$	c	e	g	b	d	f	$h \rangle$
P 's indices rotated left	=	(000	010	100	110	001	011	101	111)
$ls P$	=	$\langle a$	e	b	f	c	g	d	$h \rangle$

Figure 8.2: Permutation functions rs , ls defined in Section 8.4.1.*Rotate Index*

A class of permutation functions can be defined by the transformations on the element indices. For a powerlist of 2^n elements we associate an n -bit index with each element, where the indices are the binary representations of $0, 1, \dots, 2^n - 1$ in sequence. (For a powerlist $u \mid v$, indices for the elements in u have “0” as the highest bit and in v have “1” as the highest bit. In $u \bowtie v$, similar remarks apply for the lowest bit.) Any bijection, h , mapping indices to indices defines a permutation of the powerlist: The element with index i is moved to the position where it has index $(h i)$. Below, we consider two simple index mapping functions; the corresponding permutations of powerlists are useful in describing the shuffle-exchange network. Note that indices are not part of our theory.

A function that rotates an index to the right (by one position) has the permutation function rs (for *right shuffle*) associated with it. The definition of rs may be understood as follows. The effect of rotating an index to the right is that the lowest bit of an index becomes the highest bit; therefore, if rs is applied to $u \bowtie v$, the elements of u —those having 0 as the lowest bit—will occupy the first half of the resulting powerlist (because their indices have “0” as the highest bit, after rotation); similarly, v will occupy the second half. Analogously, the function that rotates an index to the left (by one position) induces the permutation defined by ls (for *left shuffle*), below. Figure 8.2 shows the effects of index rotations on an 8-element list.

$$\begin{aligned}
 rs\langle x \rangle &= \langle x \rangle & , & & ls\langle x \rangle &= \langle x \rangle \\
 rs(u \bowtie v) &= u \mid v & , & & ls(u \mid v) &= u \bowtie v
 \end{aligned}$$

It is trivial to see that rs , ls are inverses.

Inversion

The function inv is defined by the following function on indices. An element with index b in P has index b' in $(inv P)$, where b' is the reversal of the bit string b . Thus,

$$\begin{array}{cccccccc} & 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \\ \text{inv}\langle & a & b & c & d & e & f & g & h \rangle = \\ \langle & a & e & c & g & b & f & d & h \rangle \end{array}$$

The definition of inv is

$$\begin{aligned} \text{inv}\langle x \rangle &= \langle x \rangle \\ \text{inv}(p \mid q) &= (\text{inv } p) \bowtie (\text{inv } q) \end{aligned}$$

This function arises in a variety of contexts. In particular, inv is used to permute the output of a Fast Fourier Transform network into the correct order.

The following proof shows a typical application of structural induction.

$$\text{INV1. } \text{inv}(p \bowtie q) = (\text{inv } p) \mid (\text{inv } q)$$

Proof is by structural induction on p and q .

$$\begin{aligned} \text{Base : } & \text{inv}(\langle x \rangle \bowtie \langle y \rangle) \\ &= \{ \text{From Law L0 : } \langle x \rangle \bowtie \langle y \rangle = \langle x \rangle \mid \langle y \rangle \} \\ & \quad \text{inv}(\langle x \rangle \mid \langle y \rangle) \\ &= \{ \text{definition of inv} \} \\ & \quad \text{inv}\langle x \rangle \bowtie \text{inv}\langle y \rangle \\ &= \{ \text{inv}\langle x \rangle = \langle x \rangle, \text{inv}\langle y \rangle = \langle y \rangle. \text{ Thus, they are singletons. Applying Law L0} \} \\ & \quad \text{inv}\langle x \rangle \mid \text{inv}\langle y \rangle \end{aligned}$$

Induction :

$$\begin{aligned} & \text{inv}((r \mid s) \bowtie (u \mid v)) \\ &= \{ \text{commutativity of } \mid, \bowtie \} \\ & \quad \text{inv}((r \bowtie u) \mid (s \bowtie v)) \\ &= \{ \text{definition of inv} \} \\ & \quad \text{inv}(r \bowtie u) \bowtie \text{inv}(s \bowtie v) \\ &= \{ \text{induction} \} \\ & \quad (\text{inv } r \mid \text{inv } u) \bowtie (\text{inv } s \mid \text{inv } v) \\ &= \{ \mid, \bowtie \text{ commute} \} \\ & \quad (\text{inv } r \bowtie \text{inv } s) \mid (\text{inv } u \bowtie \text{inv } v) \\ &= \{ \text{apply definition of inv to both sides of } \mid \} \\ & \quad \text{inv}(r \mid s) \mid \text{inv}(u \mid v) \end{aligned}$$

□

Using INV1 and structural induction, it is easy to establish

$$\begin{aligned} \text{inv}(\text{inv } P) &= P, \\ \text{inv}(\text{rev } P) &= \text{rev}(\text{inv } P) \end{aligned}$$

$n = 0$ $\langle [] \rangle$
 $n = 1$ $\langle [0] [1] \rangle$
 $n = 2$ $\langle [00] [01] [11] [10] \rangle$
 $n = 3$ $\langle [000] [001] [011] [010] [110] [111] [101] [100] \rangle$

Figure 8.3: Standard Gray code sequence for n , $n = 0, 1, 2, 3$

and for any scalar operator \oplus

$$\text{inv}(P \oplus Q) = (\text{inv } P) \oplus (\text{inv } Q)$$

The last result holds for any permutation function in place of inv .

8.4.2 Reduction

In the linear list theory [5], reduction is a higher order function of two arguments, an associative binary operator and a list. Reduction applied to \oplus and $[a_0 a_1 \dots a_n]$ yields $(a_0 \oplus a_1 \oplus \dots \oplus a_n)$. This function over powerlists is defined by

$$\begin{aligned} \text{red } \oplus \langle x \rangle &= x \\ \text{red } \oplus (p \mid q) &= (\text{red } \oplus p) \oplus (\text{red } \oplus q) \end{aligned}$$

8.4.3 Gray Code

Gray code sequence [19] for n , $n \geq 0$, is a sequence of 2^n n -bit strings where the consecutive strings in the sequence differ in exactly one bit position. (The last and the first strings in the sequence are considered consecutive.) Standard Gray code sequences for $n = 0, 1, 2, 3$ are shown in Figure 8.3. We represent the n -bit strings by linear lists of length n and a Gray code sequence by a powerlist whose elements are these linear lists. The standard Gray code sequence may be computed by function G , for any n .

$$\begin{aligned} G \ 0 &= \langle [] \rangle \\ G \ (n + 1) &= (0 : P) \mid (1 : (\text{rev } P)) \\ &\text{where } P = (G \ n) \end{aligned}$$

Here, $(0 :)$ is a scalar function that takes a linear list as an argument and appends 0 as its prefix. According to the coercion rule, $(0 : P)$ is the powerlist obtained by prefixing every element of P by 0. Similarly, $(1 : (\text{rev } P))$ is defined, where the function rev is from Section 8.2.2.

8.4.4 Polynomial

A polynomial with coefficients p_j , $0 \leq j < 2^n$, where $n \geq 0$, may be represented by a powerlist p whose j^{th} element is p_j . The polynomial value at some point ω is $\sum_{0 \leq j < 2^n} p_j \times \omega^j$. For $n > 0$ this quantity is

$$\sum_{0 \leq j < 2^{n-1}} p_{2j} \times \omega^{2j} + \sum_{0 \leq j < 2^{n-1}} p_{2j+1} \times \omega^{2j+1}.$$

The following function, ep , evaluates a polynomial p using this strategy. In anticipation of the Fast Fourier Transform, we generalize ep to accept an arbitrary powerlist as its second argument. For powerlists p , w (of, possibly, unequal lengths) let $(p \text{ ep } w)$ be a powerlist of the same length as w , obtained by evaluating p at each element of w .

$$\begin{aligned} \langle x \rangle ep \ w &= \langle x \rangle \\ (p \bowtie q) \text{ ep } w &= (p \text{ ep } w^2) + (w \times (q \text{ ep } w^2)) \end{aligned}$$

Note that w^2 is the pointwise squaring of w . Also, note that ep is a pointwise function in its second argument, i.e.,

$$p \text{ ep } (u \mid v) = (p \text{ ep } u) \mid (p \text{ ep } v)$$

8.4.5 Fast Fourier Transform

For a polynomial p with complex coefficients, its Fourier transform is obtained by evaluating p at a sequence (i.e., powerlist) of points, $(W \ p)$. Here, $(W \ p)$ is the powerlist $\langle \omega^0, \omega^1, \dots, \omega^{n-1} \rangle$, where n is the length of p and ω is the n^{th} principal root of 1. Note that $(W \ p)$ depends only on the length of p but not its elements; hence, for similar powerlists p , q , $(W \ p) = (W \ q)$. It is easy to define the function W in a manner similar to ep .

We need the following properties of W for the derivation of *FFT*. Equation (1) follows from the definition of W and the fact that $\omega^{2 \times N} = 1$, where N is the length of p (and q). The second equation says that the right half of $W(p \bowtie q)$ is the negation of its left half. This is because each element in the right half is the same as the corresponding element in the left half multiplied by ω^N ; since ω is the $(2 \times N)^{\text{th}}$ root of 1, $\omega^N = -1$.

$$W^2(p \bowtie q) = (W \ p) \mid (W \ q) \tag{8.1}$$

$$W(p \bowtie q) = u \mid (-u), \text{ for some } u \tag{8.2}$$

The Fourier transform, FT , of a powerlist p is a powerlist of the same length as p , given by

$$FT \ p = p \text{ ep } (W \ p)$$

where ep is the function defined in Section 8.4.4.

The straightforward computation of $(p \text{ ep } v)$ for any p, v consists of evaluating p at each element of v ; this takes time $O(N^2)$ where p, v have length

N . Since $(W p)$ is of a special form the Fourier transform can be computed in $O(N \log N)$ steps, using the the Fast Fourier Transform algorithm [12]. This algorithm also admits an efficient parallel implementation, requiring $O(\log N)$ steps on $O(N)$ processors. We derive the FFT algorithm next.

$$\begin{aligned}
& FT\langle x \rangle \\
&= \{ \text{definition of } FT \} \\
& \quad x \text{ ep } (W\langle x \rangle) \\
&= \{ \text{Since } W\langle x \rangle \text{ is a singleton, from the definition of } ep \} \\
& \quad \langle x \rangle
\end{aligned}$$

For the general case,

$$\begin{aligned}
& FT(p \bowtie q) \\
&= \{ \text{From the definition of } FT \} \\
& \quad (p \bowtie q) \text{ ep } W(p \bowtie q) \\
&= \{ \text{from the definition of } ep \} \\
& \quad p \text{ ep } W^2(p \bowtie q) + W(p \bowtie q) \times (q \text{ ep } W^2(p \bowtie q)) \\
&= \{ \text{from the property of } W; \text{ see equation (1)} \} \\
& \quad p \text{ ep } ((W p) | (W q)) + W(p \bowtie q) \times (q \text{ ep } ((W p) | (W q))) \\
&= \{ \text{distribute each } ep \text{ over its second argument} \} \\
& \quad ((p \text{ ep } (W p)) | (p \text{ ep } (W q))) + W(p \bowtie q) \times ((q \text{ ep } (W p)) | (q \text{ ep } (W q))) \\
&= \{ (W p) = (W q), p \text{ ep } (W p) = FT p, q \text{ ep } (W q) = FT q \} \\
& \quad ((FT p) | (FT p)) + W(p \bowtie q) \times ((FT q) | (FT q)) \\
&= \{ \text{using } P, Q \text{ for } FT p, FT q, \text{ and } u | (-u) \text{ for } W(p \bowtie q); \text{ see equation (2)} \} \\
& \quad (P | P) + (u | -u) \times (Q | Q) \\
&= \{ | \text{ and } \times \text{ in the second term commute} \} \\
& \quad (P | P) + ((u \times Q) | (-u \times Q)) \\
&= \{ | \text{ and } + \text{ commute} \} \\
& \quad (P + u \times Q) | (P - u \times Q)
\end{aligned}$$

We collect the two equations for FT to define FFT , the Fast Fourier Transform. In the following, $(powers p)$ is the powerlist $\langle \omega^0, \omega^1, \dots, \omega^{N-1} \rangle$ where N is the length of p and ω is the $(2 \times N)^{th}$ principal root of 1. This was the value of u in the previous paragraph. The function $powers$ can be defined similarly to ep .

$$\begin{aligned}
& FFT\langle x \rangle = \langle x \rangle \\
& FFT(p \bowtie q) = (P + u \times Q) | (P - u \times Q) \\
& \quad \text{where} \quad P = FFT p \\
& \quad \quad \quad Q = FFT q \\
& \quad \quad \quad u = powers p
\end{aligned}$$

It is clear that $FFT(p \bowtie q)$ can be computed from $(FFT p)$ and $(FFT q)$ in $O(N)$ sequential steps or $O(1)$ parallel steps using $O(N)$ processors (u can be

computed in parallel), where N is the length of p . Therefore, $FFT(p \bowtie q)$ can be computed in $O(N \log N)$ sequential steps or, $O(\log N)$ parallel steps using $O(N)$ processors.

The compactness of this description of FFT is in striking contrast to the usual descriptions; for instance, see [10, Section 6.13]. The compactness can be attributed to the use of recursion and the avoidance of explicit indexing of the elements by employing $|$ and \bowtie . FFT illustrates the need for including both $|$ and \bowtie as constructors for powerlists. (Another function that employs both $|$ and \bowtie is `inv` of Section 8.4.1.)

Inverse Fourier Transform

The inverse of the Fourier Transform, IFT, can be defined similarly to the FFT. We derive the definition of IFT from that of the FFT by pattern matching.

For a singleton powerlist, $\langle x \rangle$, we compute

$$\begin{aligned} & IFT\langle x \rangle \\ = & \{ \langle x \rangle = FFT\langle x \rangle \} \\ & IFT(FFT\langle x \rangle) \\ = & \{ IFT, FFT \text{ are inverses} \} \\ & \langle x \rangle \end{aligned}$$

For the general case, we have to compute $IFT(r | s)$ given r, s . Let

$$IFT(r | s) = p \bowtie q$$

in the unknowns p, q . This form of deconstruction is chosen so that we can easily solve the equations we generate, next. Taking FFT of both sides,

$$FFT(IFT(r | s)) = FFT(p \bowtie q)$$

The left side is $(r | s)$ because IFT, FFT are inverses. Replacing the right side by the definition of $FFT(p \bowtie q)$ yields the following equations.

$$\begin{aligned} r | s &= (P + u \times Q) | (P - u \times Q) \\ P &= FFT\ p \\ Q &= FFT\ q \\ u &= \text{powers } p \end{aligned}$$

These equations are easily solved for the unknowns P, Q, u, p, q . (The law of unique deconstruction, L2, can be used to deduce from the first equation that $r = P + u \times Q$ and $s = P - u \times Q$. Also, since p and r are of the same length we may define u using r instead of p .) The solutions of these equations yield the following definition for IFT. Here, $/2$ divides each element of the given powerlist by 2.

$$\begin{aligned} IFT\langle x \rangle &= \langle x \rangle \\ IFT(r | s) &= p \bowtie q \\ \text{where } P &= (r + s)/2 \end{aligned}$$

$$\begin{aligned}
 u &= \text{powers } r \\
 Q &= ((r - s)/2)/u \\
 p &= IFT P \\
 q &= IFT Q
 \end{aligned}$$

As in the FFT, the definition of IFT includes both constructors, $|$ and \bowtie . It can be implemented efficiently on a butterfly network. The complexity of IFT is same as that of the FFT.

8.4.6 Batcher Sort

In this section, we develop some elementary results about sorting and discuss two remarkable sorting methods due to Batcher[4]. We find it interesting that \bowtie (not $|$) is the preferred operator in discussing the principles of parallel sorting. Henceforth, a list is *sorted* means that its elements are arranged in non-decreasing order.

A general method of sorting is given by

$$\begin{aligned}
 \text{sort}\langle x \rangle &= \langle x \rangle \\
 \text{sort}(p \bowtie q) &= (\text{sort } p) \text{ merge } (\text{sort } q)
 \end{aligned}$$

where *merge* (written as a binary infix operator) creates a single sorted powerlist out of the elements of its two argument powerlists each of which is sorted. In this section, we show two different methods for implementing *merge*. One scheme is Batcher merge, given by the operator *bm*. Another scheme is given by *bitonic* sort where the sorted lists u, v are merged by applying the function *bi* to $(u | (\text{rev } v))$.

A comparison operator, \updownarrow , is used in these algorithms. The operator is applied to a pair of equal length powerlists, p, q ; it creates a single powerlist out of the elements of p, q by

$$p \updownarrow q = (p \text{ min } q) \bowtie (p \text{ max } q)$$

That is, the $2i^{\text{th}}$ and $(2i + 1)^{\text{th}}$ items of $p \updownarrow q$ are $(p_i \text{ min } q_i)$ and $(p_i \text{ max } q_i)$, respectively. The powerlist $p \updownarrow q$ can be computed in constant time using $O(\text{len } p)$ processors.

Bitonic Sort

A sequence of numbers, $x_0, x_1, \dots, x_i, \dots, x_N$, is *bitonic* if there is an index i , $0 \leq i \leq N$, such that x_0, x_1, \dots, x_i is monotonic (ascending or descending) and x_i, \dots, x_N is monotonic. The function *bi*, given below, applied to a bitonic powerlist returns a sorted powerlist of the original items.

$$\begin{aligned}
 \text{bi}\langle x \rangle &= \langle x \rangle \\
 \text{bi}(p \bowtie q) &= (\text{bi } p) \updownarrow (\text{bi } q)
 \end{aligned}$$

For sorted powerlists u, v , the powerlist $(u \mid (\text{rev } v))$ is bitonic; thus u, v can be merged by applying bi to $(u \mid (\text{rev } v))$. The form of the recursive definition suggests that bi can be implemented on $O(N)$ processors in $O(\log N)$ parallel steps, where N is the length of the argument powerlist.

Batcher Merge

Batcher has also proposed a scheme for merging two sorted lists. We define this scheme, bm , as an infix operator below.

$$\begin{aligned} \langle x \rangle bm \langle y \rangle &= \langle x \rangle \downarrow \langle y \rangle \\ (r \bowtie s) bm (u \bowtie v) &= (r bm v) \downarrow (s bm u) \end{aligned}$$

The function bm is well-suited for parallel implementation. The recursive form suggests that $(r bm v)$ and $(s bm u)$ can be computed in parallel. Since \downarrow can be applied in $O(1)$ parallel steps using $O(N)$ processors, where N is the length of the argument powerlists, the function bm can be evaluated in $O(\log N)$ parallel steps. In the rest of this section, we develop certain elementary facts about sorting and prove the correctness of bi and bm .

Elementary Facts about Sorting

We consider only “compare and swap” type sorting methods. It is known (see [29]) that such a sorting scheme is correct if and only if it sorts lists containing 0’s and 1’s only. Therefore, we restrict our discussion to powerlists containing 0’s and 1’s, only.

For a powerlist p , let $(z p)$ be the number of 0’s in it. To simplify notation, we omit the space and write zp . Clearly,

$$A0. \quad z(p \bowtie q) = zp + zq \text{ and } z\langle x \rangle \text{ is either } 0 \text{ or } 1.$$

Powerlists containing only 0’s and 1’s have the following properties.

- A1. $\langle x \rangle$ sorted and $\langle x \rangle$ bitonic.
- A2. $(p \bowtie q)$ sorted $\equiv p$ sorted $\wedge q$ sorted $\wedge 0 \leq zp - zq \leq 1$
- A3. $(p \bowtie q)$ bitonic $\Rightarrow p$ bitonic $\wedge q$ bitonic $\wedge |zp - zq| \leq 1$

Note : The condition analogous to (A2) under which $p \mid q$ is sorted is,

$$A2'. \quad (p \mid q) \text{ sorted} \equiv p \text{ sorted} \wedge q \text{ sorted} \wedge (zp < (\text{len } p) \Rightarrow zq = 0)$$

The simplicity of (A2), compared with (A2'), may suggest why \bowtie is the primary operator in parallel sorting. \square

The following results, (B1, B2), are easy to prove. We prove (B3).

- B1. p sorted, q sorted, $zp \geq zq \Rightarrow (p \min q) = p \wedge (p \max q) = q$ \square
- B2. $z(p \downarrow q) = zp + zq$ \square

B3. p sorted, q sorted, $|zp - zq| \leq 1 \Rightarrow (p \uparrow q)$ sorted

Proof: Since the statement of B3 is symmetric in p, q , assume $zp \geq zq$.

$$\begin{aligned}
 & p \text{ sorted, } q \text{ sorted, } |zp - zq| \leq 1 \\
 \Rightarrow & \quad \{\text{assumption: } zp \geq zq\} \\
 & p \text{ sorted, } q \text{ sorted, } 0 \leq zp - zq \leq 1 \\
 \Rightarrow & \quad \{\text{A2 and B1}\} \\
 & p \bowtie q \text{ sorted, } (p \min q) = p, (p \max q) = q \\
 \Rightarrow & \quad \{\text{replace } p, q \text{ in } p \bowtie q \text{ by } (p \min q), (p \max q)\} \\
 & (p \min q) \bowtie (p \max q) \text{ sorted} \\
 \Rightarrow & \quad \{\text{definition of } p \uparrow q\} \\
 & p \uparrow q \text{ sorted}
 \end{aligned}$$

Correctness of Bitonic Sort

We show that the function bi applied to a bitonic powerlist returns a sorted powerlist of the original elements: (B4) states that bi preserves the number of zeroes of its argument list (i.e., it loses no data) and (B5) states that the resulting list is sorted.

B4. $z(bi\ p) = zp$

Proof: By structural induction, using B2. □

B5. L bitonic $\Rightarrow (bi\ L)$ sorted

Proof: By structural induction.

Base: Straightforward.

Induction: Let $L = p \bowtie q$

$$\begin{aligned}
 & p \bowtie q \text{ bitonic} \\
 \Rightarrow & \quad \{\text{A3}\} \\
 & p \text{ bitonic, } q \text{ bitonic, } |zp - zq| \leq 1 \\
 \Rightarrow & \quad \{\text{induction on } p \text{ and } q\} \\
 & (bi\ p) \text{ sorted, } (bi\ q) \text{ sorted, } |zp - zq| \leq 1 \\
 \Rightarrow & \quad \{\text{from B4: } z(bi\ p) = zp, z(bi\ q) = zq\} \\
 & (bi\ p) \text{ sorted, } (bi\ q) \text{ sorted, } |z(bi\ p) - z(bi\ q)| \leq 1 \\
 \Rightarrow & \quad \{\text{apply B3 with } (bi\ p), (bi\ q) \text{ for } p, q\} \\
 & (bi\ p) \uparrow (bi\ q) \text{ sorted} \\
 \Rightarrow & \quad \{\text{definition of } bi\} \\
 & bi(p \bowtie q) \text{ sorted}
 \end{aligned}$$

Correctness of Batcher Merge

We can show that bm merges two sorted powerlists in a manner similar to the proof of bi . Instead, we establish a simple relationship between the functions bm and bi from which the correctness of the former is obvious. We show that

B6. $p \text{ } bm \text{ } q = bi(p \mid (rev \text{ } q))$, where rev reverses a powerlist (Section 8.2.2).

If p, q are sorted then $p \mid (rev \text{ } q)$ is bitonic (a fact that we don't prove here). Then, from the correctness of bi it follows that $bi(p \mid (rev \text{ } q))$ and, hence, $p \text{ } bm \text{ } q$ is sorted (and it contains the elements of p and q).

Proof of B6: By structural induction.

Base: Let $p, q = \langle x \rangle, \langle y \rangle$

$$\begin{aligned}
 & bi(\langle x \rangle \mid rev\langle y \rangle) \\
 = & \{\text{definition of } rev\} \\
 & bi(\langle x \rangle \mid \langle y \rangle) \\
 = & \{(\langle x \rangle \mid \langle y \rangle) = (\langle x \rangle \bowtie \langle y \rangle)\} \\
 & bi(\langle x \rangle \bowtie \langle y \rangle) \\
 = & \{\text{definition of } bi\} \\
 & \langle x \rangle \updownarrow \langle y \rangle \\
 = & \{\text{definition of } bm\} \\
 & \langle x \rangle \text{ } bm \text{ } \langle y \rangle
 \end{aligned}$$

Induction: Let $p, q = r \bowtie s, u \bowtie v$

$$\begin{aligned}
 & bi(p \mid (rev \text{ } q)) \\
 = & \{\text{expanding } p, q\} \\
 & bi((r \bowtie s) \mid rev(u \bowtie v)) \\
 = & \{\text{definition of } rev\} \\
 & bi((r \bowtie s) \mid (rev \text{ } v \bowtie rev \text{ } u)) \\
 = & \{ \mid, \bowtie \text{ commute} \} \\
 & bi((r \mid rev \text{ } v) \bowtie (s \mid rev \text{ } u)) \\
 = & \{\text{definition of } bi\} \\
 & bi(r \mid rev \text{ } v) \updownarrow bi(s \mid rev \text{ } u) \\
 = & \{\text{induction}\} \\
 & (r \text{ } bm \text{ } v) \updownarrow (s \text{ } bm \text{ } u) \\
 = & \{\text{definition of } bm\} \\
 & (r \bowtie s) \text{ } bm \text{ } (u \bowtie v)
 \end{aligned}$$

= {using the definitions of p, q }
 $p \text{ } b m \text{ } q$ □

The compactness of the description of Batcher's sorting schemes and the simplicity of their correctness proofs demonstrate the importance of treating recursion and parallelism simultaneously.

8.4.7 Prefix Sum

Let L be a powerlist of scalars and \oplus be a binary, associative operator on that scalar type. The prefix sum of L with respect to \oplus , ($ps \ L$), is a list of the same length as L given by

$$ps \langle x_0, x_1, \dots, x_i, \dots, x_N \rangle = \langle x_0, x_0 \oplus x_1, \dots, x_0 \oplus x_1 \oplus \dots \oplus x_i, \dots, x_0 \oplus x_1 \oplus \dots \oplus x_N \rangle,$$

that is, in ($ps \ L$) the element with index i , $i > 0$, is obtained by applying \oplus to the first $(i + 1)$ elements of L in order. We will give a formal definition of prefix sum later in this section.

Prefix sum is of fundamental importance in parallel computing. We show that two known algorithms for this problem can be concisely represented and proved in our theory. Again, zip turns out to be the primary operator for describing these algorithms.

A particularly simple scheme for prefix sum of 8 elements is shown in Figure 8.4. In that figure, the numbered nodes represent processors, though the same 8 physical processors are used at all levels. Initially, processor i holds the list element L_i , for all i . The connections among the processors at different levels depict data transmissions. In level 0, each processor, from 0 through 6, sends its data to its right neighbor. In the i^{th} level, processor i sends its data to $(i + 2^i)$, if such a processor exists (this means that for $j < 2^i$, processor j receives no data in level i data transmission). Each processor updates its own data, d , to $r \oplus d$ where r is the data it receives; if it receives no data in some level then d is unchanged. It can be shown that after completion of the computation at level $(\log_2(\text{len } L))$, processor i holds the i^{th} element of ($ps \ L$).

Another scheme, due to Ladner and Fischer[32], first applies \oplus to adjacent elements x_{2i}, x_{2i+1} to compute the list $\langle x_0 \oplus x_1, \dots, x_{2i} \oplus x_{2i+1}, \dots \rangle$. This list has half as many elements as the original list; its prefix sum is then computed recursively. The resulting list is $\langle x_0 \oplus x_1, \dots, x_0 \oplus x_1 \oplus \dots \oplus x_{2i} \oplus x_{2i+1}, \dots \rangle$. This list contains half of the elements of the final list; the missing elements are $x_0, x_0 \oplus x_1 \oplus x_2, \dots, x_0 \oplus x_1 \oplus \dots \oplus x_{2i}, \dots$. These elements can be computed by "adding" x_2, x_4, \dots , appropriately to the elements of the already computed list.

Both schemes for prefix computation are inherently recursive. Our formulations will highlight both parallelism and recursion.

Specification

As we did for the sorting schemes (Section 8.4.6), we introduce an operator in terms of which the prefix sum problem can be defined. First, we postulate




Figure 8.4: A network to compute the prefix sum of 8 elements.

that 0 is the left identity element of \oplus , i.e., $0 \oplus x = x$. For a powerlist p , let p^* be the powerlist obtained by shifting p to the right by one. The effect of shifting is to append a 0 to the left and discard the rightmost element of p ; thus, $\langle a b c d \rangle^* = \langle 0 a b c \rangle$. Formally,

$$\begin{aligned} \langle x \rangle^* &= \langle 0 \rangle \\ (p \bowtie q)^* &= q^* \bowtie p \end{aligned}$$

It is easy to show

- S1. $(r \oplus s)^* = r^* \oplus s^*$
- S2. $(p \bowtie q)^{**} = p^* \bowtie q^*$

Consider the following equation in the powerlist variable z .

$$z = z^* \oplus L \tag{DE}$$

where L is some given powerlist. This equation has a unique solution in z , because

$$\begin{aligned} z_0 &= (z^*)_0 \oplus L_0 \\ &= 0 \oplus L_0 \\ &= L_0 \quad , \text{ and} \\ z_{i+1} &= z_i \oplus L_{i+1} \quad , \quad 0 \leq i < (\text{len } L) - 1 \end{aligned}$$

For $L = \langle a b c d \rangle$, $z = \langle a \ a \oplus b \ a \oplus b \oplus c \ a \oplus b \oplus c \oplus d \rangle$ which is exactly $(ps L)$. We define $(ps L)$ to be the unique solution of (DE), and we call (DE) the defining equation for $(ps L)$.

Notes

1. The operator \oplus is not necessarily commutative. Therefore, the rhs of (DE) may not be the same as $L \oplus z^*$.

2. The operator \oplus is scalar; so, it commutes with \bowtie .
3. The uniqueness of the solution of (DE) can be proved entirely within the powerlist algebra, similar to the derivation of Ladner-Fischer scheme given later in this section.
4. Adams[1] has specified the prefix-sum problem without postulating an explicit “0” element. For any \oplus , he introduces a binary operator $\vec{\oplus}$ over two similar powerlists such that $p\vec{\oplus}q = p^* \oplus q$. The operator $\vec{\oplus}$ can be defined without introducing a “0”.

Computation of the Prefix Sum

The function *sps* (simple prefix sum) defines the scheme of Figure 8.4.

$$\begin{aligned} \text{sps } \langle x \rangle &= \langle x \rangle \\ \text{sps } L &= (\text{sps } u) \bowtie (\text{sps } v) \\ &\text{where } u \bowtie v = L^* \oplus L \end{aligned}$$

In the first level in Figure 8.4, $L^* \oplus L$ is computed. If $L = \langle x_0, x_1, \dots, x_i, \dots \rangle$ then this is $\langle x_0, x_0 \oplus x_1, \dots, x_i \oplus x_{i+1}, \dots \rangle$. This is the zip of the two sublists $\langle x_0, x_1 \oplus x_2, \dots, x_{2i-1} \oplus x_{2i}, \dots \rangle$ and $\langle x_0 \oplus x_1, \dots, x_{2i} \oplus x_{2i+1}, \dots \rangle$. Next, prefix sums of these two lists are computed (independently) and then zipped.

The Ladner-Fischer scheme is defined by the function *lf*.

$$\begin{aligned} \text{lf } \langle x \rangle &= \langle x \rangle \\ \text{lf } (p \bowtie q) &= (t^* \oplus p) \bowtie t \\ &\text{where } t = \text{lf } (p \oplus q) \end{aligned}$$

Correctness

We can prove the correctness of *sps* and *lf* by showing that the function *ps* satisfies the equations defining each of these functions. It is more instructive to see that both *sps* and *lf* can be derived easily from the specification (DE). We carry out this derivation for the Fischer-Ladner scheme as an illustration of the power of algebraic manipulations. First, we note, $ps \langle x \rangle = \langle x \rangle$.

$$\begin{aligned} &ps \langle x \rangle \\ = &\quad \{\text{from the defining equation DE for } ps \langle x \rangle\} \\ &(ps \langle x \rangle)^* \oplus \langle x \rangle \\ = &\quad \{\text{definition of } *\} \\ &\langle 0 \rangle \oplus \langle x \rangle \\ = &\quad \{\oplus \text{ is a scalar operation}\} \\ &\langle 0 \oplus x \rangle \\ = &\quad \{0 \text{ is the identity of } \oplus\} \\ &\langle x \rangle \end{aligned}$$

Derivation of Ladner-Fischer Scheme

Given a powerlist $p \bowtie q$, we derive an expression for $ps(p \bowtie q)$. Let $r \bowtie t$, in unknowns r, t , be $ps(p \bowtie q)$. We solve for r, t .

$$\begin{aligned}
 & r \bowtie t \\
 = & \quad \{r \bowtie t = ps(p \bowtie q). \text{ Using (DE)}\} \\
 & (r \bowtie t)^* \oplus (p \bowtie q) \\
 = & \quad \{(r \bowtie t)^* = t^* \bowtie r\} \\
 & (t^* \bowtie r) \oplus (p \bowtie q) \\
 = & \quad \{\oplus, \bowtie \text{ commute}\} \\
 & (t^* \oplus p) \bowtie (r \oplus q)
 \end{aligned}$$

Applying law L2 (unique deconstruction) to the equation $r \bowtie t = (t^* \oplus p) \bowtie (r \oplus q)$, we conclude that

$$LF1. \quad r = t^* \oplus p, \text{ and}$$

$$LF2. \quad t = r \oplus q$$

Now, we eliminate r from (LF2) using (LF1) to get $t = t^* \oplus p \oplus q$. Using (DE) and this equation we obtain

$$LF3. \quad t = ps(p \oplus q)$$

We summarize the derivation of $ps(p \bowtie q)$.

$$\begin{aligned}
 & ps(p \bowtie q) \\
 = & \quad \{\text{by definition}\} \\
 & r \bowtie t \\
 = & \quad \{\text{Using (LF1) for } r\} \\
 & (t^* \oplus p) \bowtie t
 \end{aligned}$$

where t is defined by LF3. This is exactly the definition of the function lf for a non-singleton powerlist. We also note that

$$\begin{aligned}
 & r \\
 = & \quad \{\text{by eliminating } t \text{ from (LF1) using (LF2)}\} \\
 & (r \oplus q)^* \oplus p \\
 = & \quad \{\text{definition of } *\} \\
 & r^* \oplus q^* \oplus p
 \end{aligned}$$

Using (DE) and this equation we obtain LF4 that is used in proving the correctness of *sps*, next.

$$LF4. \quad r = ps(q^* \oplus p)$$

Correctness of *sps*

We show that for a non-singleton powerlist L ,

$$ps L = (ps u) \bowtie (ps v), \text{ where } u \bowtie v = L^* \oplus L.$$

Proof: Let $L = p \bowtie q$. Then

$$\begin{aligned} & ps L \\ = & \{L = p \bowtie q\} \\ & ps(p \bowtie q) \\ = & \{ps(p \bowtie q) = r \bowtie t, \text{ where } r, t \text{ are given by (LF4,LF3)}\} \\ & ps(q^* \oplus p) \bowtie ps(p \oplus q) \\ = & \{\text{Letting } u = q^* \oplus p, v = p \oplus q\} \\ & (ps u) \bowtie (ps v) \end{aligned}$$

Now, we show that $u \bowtie v = L^* \oplus L$.

$$\begin{aligned} & u \bowtie v \\ = & \{u = q^* \oplus p, v = p \oplus q\} \\ & (q^* \oplus p) \bowtie (p \oplus q) \\ = & \{\oplus, \bowtie \text{ commute}\} \\ & (q^* \bowtie p) \oplus (p \bowtie q) \\ = & \{\text{Apply the definition of } * \text{ to the first term}\} \\ & (p \bowtie q)^* \oplus (p \bowtie q) \\ = & \{L = p \bowtie q\} \\ & L^* \oplus L \end{aligned}$$

Remarks. A more traditional way of describing a prefix sum algorithm, such as the simple scheme of Figure 8.4, is to explicitly name the quantities that are being computed, and establish relationships among them. Let y_{ij} be computed by the i^{th} processor at the j^{th} level. Then, for all i, j , $0 \leq i < 2^n$, $0 \leq j < n$, where n is the logarithmic length of the list,

$$y_{i0} = x_i, \text{ and} \\ y_{i,j+1} = \left\{ \begin{array}{ll} y_{i-2^j,j} & , \quad i \geq 2^j \\ 0 & , \quad i < 2^j \end{array} \right\} \oplus y_{ij}$$

The correctness criterion is

$$y_{in} = x_0 \oplus .. \oplus x_i$$

This description is considerably more difficult to manipulate. The parallelism in it is harder to see. The proof of correctness requires manipulations of indices: for this example, we have to show that for all i, j

$$y_{ij} = x_k \oplus .. \oplus x_i$$

where $k = \max(0, i - 2^j + 1)$.

The Ladner-Fischer scheme is even more difficult to specify in this manner. Algebraic methods are to be preferred for describing uniform operations on aggregates of data.

8.5 Higher Dimensional Arrays

A major part of parallel computing involves arrays of one or more dimensions. An array of m dimensions (dimensions are numbered 0 through $m - 1$) is represented by a powerlist of depth $(m - 1)$. Conversely, since powerlist elements are similar, a powerlist of depth $(m - 1)$ may be regarded as an array of dimension m . For instance, a matrix of r rows and c columns may be represented as a powerlist of c elements, each element being a powerlist of length r storing the items of a column; conversely, the same matrix may be represented by a powerlist of r elements, each element being a powerlist of c elements.

In manipulating higher dimensional arrays we prefer to think in terms of array operations rather than operations on nested powerlists. Therefore, we introduce construction operators, analogous to $|$ and \bowtie , for tie and zip along any specified dimension. We use $|'$, \bowtie' for the corresponding operators in dimension 1, $|''$, \bowtie'' for the dimension 2, etc. The definitions of these operators are in Section 8.5.2; for the moment it is sufficient to regard $|'$ as the point-wise application of $|$ to the argument powerlists (and similarly, \bowtie'). Thus, for similar (power) matrices A, B that are stored columnwise (i.e., each element is a column), $A | B$ is the concatenation of A, B by rows and $A |' B$ is their concatenation by columns. Figure 8.5 shows applications of these operators on specific matrices.

Given these constructors we may define a matrix to be either

- a singleton matrix $\langle\langle x \rangle\rangle$, or
- $p | q$ where p, q are (similar) matrices, or
- $u |' v$ where u, v are (similar) matrices.

Analogous definitions can be given for n -dimensional arrays. Observe that the length of each dimension is a power of 2. As we had in the case of a powerlist, the same matrix can be constructed in several different ways, say, first

$$\begin{array}{ccc}
 A = \left\langle \begin{array}{c} \wedge \quad \wedge \\ 2 \quad 4 \\ 3 \quad 5 \\ \vee \quad \vee \end{array} \right\rangle & & B = \left\langle \begin{array}{c} \wedge \quad \wedge \\ 0 \quad 1 \\ 6 \quad 7 \\ \vee \quad \vee \end{array} \right\rangle \\
 \\
 A | B = \left\langle \begin{array}{c} \wedge \quad \wedge \quad \wedge \quad \wedge \\ 2 \quad 4 \quad 0 \quad 1 \\ 3 \quad 5 \quad 6 \quad 7 \\ \vee \quad \vee \quad \vee \quad \vee \end{array} \right\rangle & & A \bowtie B = \left\langle \begin{array}{c} \wedge \quad \wedge \quad \wedge \quad \wedge \\ 2 \quad 0 \quad 4 \quad 1 \\ 3 \quad 6 \quad 5 \quad 7 \\ \vee \quad \vee \quad \vee \quad \vee \end{array} \right\rangle \\
 \\
 A |' B = \left\langle \begin{array}{c} \wedge \quad \wedge \\ 2 \quad 4 \\ 3 \quad 5 \\ 0 \quad 1 \\ 6 \quad 7 \\ \vee \quad \vee \end{array} \right\rangle & & A \bowtie' B = \left\langle \begin{array}{c} \wedge \quad \wedge \\ 2 \quad 4 \\ 0 \quad 1 \\ 3 \quad 5 \\ 6 \quad 7 \\ \vee \quad \vee \end{array} \right\rangle
 \end{array}$$

Figure 8.5: Applying $|$, \bowtie , $|'$, \bowtie' over matrices. Matrices are stored by columns. Typical matrix format is used for display, though each matrix is to be regarded as a powerlist of powerlists.

by constructing the rows and then the columns, or vice versa. We will show, in Section 8.5.2, that

$$(p | q) |' (u | v) = (p |' u) | (q |' v)$$

i.e., $|$, $|'$ commute.

Note : We could have defined a matrix using \bowtie and \bowtie' instead of $|$ and $|'$. As $|$ and \bowtie are duals in the sense that either can be used to construct (or uniquely deconstruct) a powerlist, $|'$ and \bowtie' are also duals, as we show in Section 8.5.2. Therefore, we will freely use all four construction operators for matrices. \square

Example : (Matrix Transposition)

Let τ be a function that transposes matrices. From the definition of a matrix, we have to consider three cases in defining τ .

$$\begin{aligned}
 \tau\langle\langle x \rangle\rangle &= \langle\langle x \rangle\rangle \\
 \tau(p | q) &= (\tau p) |' (\tau q) \\
 \tau(u |' v) &= (\tau u) | (\tau v)
 \end{aligned}$$

The description of function τ , though straightforward, has introduced the possibility of an inconsistent definition. For a 2×2 matrix, for instance, either of the last two deconstructions apply, and it is not obvious that the same result is obtained independent of the order in which the rules are applied. We show that τ is a function.

We prove the result by structural induction. For a matrix of the form $\langle\langle x \rangle\rangle$, only the first deconstruction applies, and, hence, the claim holds. Next, consider

$$\sigma \begin{array}{|c|c|} \hline p & q \\ \hline u & v \\ \hline \end{array} = \begin{array}{|c|c|} \hline \sigma p & \sigma u \\ \hline \sigma q & \sigma v \\ \hline \end{array}$$

Figure 8.6: Schematic of the transposition of a square powermatrix.

a matrix to which both of the last two deconstructions apply. Such a matrix is of the form $(p \mid q) \mid' (u \mid v)$ which, as remarked above, is also $(p \mid' u) \mid (q \mid' v)$. Applying one step of each of the last two rules in different order, we get

$$\begin{aligned} & \tau((p \mid q) \mid' (u \mid v)) \\ &= \{ \text{applying the last rule} \} \\ & \quad ((\tau p \mid q) \mid (\tau u \mid v)) \\ &= \{ \text{applying the middle rule} \} \\ & \quad ((\tau p) \mid' (\tau q)) \mid ((\tau u) \mid' (\tau v)) \end{aligned}$$

And,

$$\begin{aligned} & \tau((p \mid' u) \mid (q \mid' v)) \\ &= \{ \text{applying first the middle rule, then the last rule} \} \\ & \quad ((\tau p) \mid (\tau u)) \mid' ((\tau q) \mid (\tau v)) \\ &= \{ \mid, \mid' \text{ commute} \} \\ & \quad ((\tau p) \mid' (\tau q)) \mid ((\tau u) \mid' (\tau v)) \end{aligned}$$

From the induction hypothesis, (τp) , (τq) , etc., are well defined. Hence,

$$\tau((p \mid q) \mid' (u \mid v)) = \tau((p \mid' u) \mid (q \mid' v))$$

Crucial to the above proof is the fact that \mid and \mid' commute; this is reminiscent of the ‘‘Church-Rosser Property’’ [11] in term rewriting systems. Commutativity is so important that we discuss it further in the next subsection.

It is easy to show that

$$\begin{aligned} \tau(p \bowtie q) &= (\tau p) \bowtie' (\tau q) \text{ and} \\ \tau(u \bowtie' v) &= (\tau u) \bowtie (\tau v) \end{aligned}$$

Transposition of a square (power) matrix can be defined by deconstructing the matrix into quarters, transposing them individually and rearranging them, as shown in Figure 8.6. From the transposition function τ for general matrices, we get a function σ for transpositions of square matrices

$$\begin{aligned} \sigma \langle \langle x \rangle \rangle &= \langle \langle x \rangle \rangle \\ \sigma((p \mid q) \mid' (u \mid v)) &= ((\sigma p) \mid' (\sigma q)) \mid ((\sigma u) \mid' (\sigma v)) \end{aligned}$$

Note the effectiveness of pattern matching in this definition.

8.5.1 Pointwise Application

Let g be a function mapping items of type α to type β . Then g' maps a powerlist of α -items to a powerlist of β -items.

$$\begin{aligned} g'\langle x \rangle &= \langle g x \rangle \\ g'(r \mid s) &= (g' r) \mid (g' s) \end{aligned}$$

Similarly, for a binary operator op

$$\begin{aligned} \langle x \rangle op' \langle y \rangle &= \langle x op y \rangle \\ (r \mid s) op' (u \mid v) &= (r op' u) \mid (s op' v) \end{aligned}$$

We have defined these two forms explicitly because we use one or the other in all our examples; f' for a function f of arbitrary arity is similarly defined. Observe that f' applied to a powerlist of length N yields a powerlist of length N . The number of primes over f determines the dimension at which f is applied (the outermost dimension is numbered 0; therefore writing \bowtie , for instance, without primes, simply zips two lists). The operator for pointwise application also appears in [3] and in [49].

Common special cases for the binary operator, op , are \mid and \bowtie and their pointwise application operators. In particular, writing \bowtie^m to denote $\bowtie \overbrace{\cdots}^m$, we define, $\bowtie^0 = \bowtie$ and for $m > 0$,

$$\begin{aligned} \langle x \rangle \bowtie^m \langle y \rangle &= \langle x \bowtie^{m-1} y \rangle \\ (r \mid s) \bowtie^m (u \mid v) &= (r \bowtie^m u) \mid (s \bowtie^m v) \end{aligned}$$

From the definition of f' , we conclude that f' and \mid commute. Below, we prove that f' commutes with \bowtie .

Theorem 1 f', \bowtie commute.

Proof: We prove the result for unary f ; the general case is similar. Proof is by structural induction.

$$\begin{aligned} \text{Base: } & f'(\langle x \rangle \bowtie \langle y \rangle) \\ &= \{ \langle x \rangle \bowtie \langle y \rangle = \langle x \rangle \mid \langle y \rangle \} \\ & \quad f'(\langle x \rangle \mid \langle y \rangle) \\ &= \{ \text{definition of } f' \} \\ & \quad f'\langle x \rangle \mid f'\langle y \rangle \\ &= \{ f'\langle x \rangle, f'\langle y \rangle = \langle f x \rangle, \langle f y \rangle. \text{ These are singleton lists} \} \\ & \quad f'\langle x \rangle \bowtie f'\langle y \rangle \end{aligned}$$

Induction:

$$\begin{aligned} & f'((p \mid q) \bowtie (u \mid v)) \\ &= \{ \mid, \bowtie \text{ in the argument commute} \} \end{aligned}$$

$$\begin{aligned}
& f'((p \bowtie u) | (q \bowtie v)) \\
= & \{f', | \text{ commute}\} \\
& f'(p \bowtie u) | f'(q \bowtie v) \\
= & \{\text{induction}\} \\
& ((f' p) \bowtie (f' u)) | ((f' q) \bowtie (f' v)) \\
= & \{|, \bowtie \text{ commute}\} \\
& ((f' p) | (f' q)) \bowtie ((f' u) | (f' v)) \\
= & \{f', | \text{ commute}\} \\
& (f'(p | q)) \bowtie (f'(u | v)) \quad \square
\end{aligned}$$

Theorem 2 For a scalar function f , $f' = f$.

Proof: Proof by structural induction is straightforward. □

Theorem 3 If f, g commute then so do f', g' .

Proof: By structural induction. □

The following results about commutativity can be derived from Theorems 1,2,3. In the following, m, n are natural numbers.

- C1. For any f and $m > n$, $f^m, |^n$ commute, and f^m, \bowtie^n commute.
- C2. For $m \neq n$, $|^m, |^n$ commute, and \bowtie^m, \bowtie^n commute.
- C3. For all m, n , $|^m, \bowtie^n$ commute.
- C4. For any scalar function f , $f, |^m$ commute, and f, \bowtie^n commute.

C1 follows by applying induction on Theorems 1 and 3 (and the fact that $f', |$ commute). C2 follows from C1; C3 from C1, Law L3 and Theorem 3; C4 from C1 and Theorem 2.

8.5.2 Deconstruction

In this section we show that any powerlist that can be written as $p |^m q$ for some p, q can also be written as $u \bowtie^m v$ for some u, v and vice versa; this is analogous to Law L1, for dual deconstruction. Analogous to Law L2, we show that such deconstructions are unique.

Theorem 4 (*dual deconstruction*): For any p, q and $m \geq 0$, if $p |^m q$ is defined then there exist u, v such that

$$u \bowtie^m v = p |^m q$$

Conversely, for any u, v and $m \geq 0$, if $u \bowtie^m v$ is defined then there exist some p, q such that

$$p |^m q = u \bowtie^m v \quad \square$$

We do not prove this theorem; its proof is similar to the theorem given below.

Theorem 5 (*unique deconstruction*): Let \otimes be $|$ or \bowtie . For any natural number m ,

$$(p \otimes^m q = u \otimes^m v) \equiv (p = u \wedge q = v)$$

Proof: Proof is by induction on m .

$m = 0$: The result follows from Law L2.

$m = n + 1$: Assume that $\otimes = |$. The proof is similar for $\otimes = \bowtie$. We prove the result by structural induction on p .

$$\begin{aligned} \text{Base: } & p = \langle a \rangle, q = \langle b \rangle, u = \langle c \rangle, v = \langle d \rangle \\ & \langle a \rangle |^{n+1} \langle b \rangle = \langle c \rangle |^{n+1} \langle d \rangle \\ & \equiv \{\text{definition of } |^{n+1}\} \\ & \langle a \rangle |^n b = \langle c \rangle |^n d \\ & \equiv \{\text{unique deconstruction using Law L2}\} \\ & a |^n b = c |^n d \\ & \equiv \{\text{induction on } n\} \\ & (a = c) \wedge (b = d) \\ & \equiv \{\text{Law L2}\} \\ & (\langle a \rangle = \langle c \rangle) \wedge (\langle b \rangle = \langle d \rangle) \end{aligned}$$

$$\begin{aligned} \text{Induction: } & p = p_0 | p_1, q = q_0 | q_1, u = u_0 | u_1, v = v_0 | v_1 \\ & (p_0 | p_1) |^{n+1} (q_0 | q_1) = (u_0 | u_1) |^{n+1} (v_0 | v_1) \\ & \equiv \{\text{definition of } |^{n+1}\} \\ & (p_0 |^{n+1} q_0) | (p_1 |^{n+1} q_1) = (u_0 |^{n+1} v_0) | (u_1 |^{n+1} v_1) \\ & \equiv \{\text{unique deconstruction using Law L2}\} \\ & (p_0 |^{n+1} q_0) = (u_0 |^{n+1} v_0) \wedge (p_1 |^{n+1} q_1) = (u_1 |^{n+1} v_1) \\ & \equiv \{\text{induction on the length of } p_0, q_0, p_1, q_1\} \\ & (p_0 = u_0) \wedge (q_0 = v_0) \wedge (p_1 = u_1) \wedge (q_1 = v_1) \\ & \equiv \{\text{Law L2}\} \\ & (p_0 | p_1) = (u_0 | u_1) \wedge (q_0 | q_1) = (v_0 | v_1) \end{aligned}$$

Theorems 4 and 5 allow a richer variety of pattern matching in function definitions, as we did for matrix transposition. We may employ $|^m, \bowtie^n$ for any natural m, n to construct a pattern over which a function can be defined.

8.5.3 Embedding Arrays in Hypercubes

An n -dimensional *hypercube* is a graph of 2^n nodes, $n \geq 0$, where each node has a unique n -bit label. Two nodes are *neighbors*, i.e., there is an edge between them, exactly when their labels differ in a single bit. Therefore, every node has n neighbors. We may represent a n -dimensional hypercube as a powerlist of depth n ; each level, except the innermost, consists of two powerlists. The operators $|^m, \bowtie^n$ for natural m, n can be used to access the nodes in any one (or any combination of) dimensions.

We conclude with an example that shows how higher dimensional structures, such as hypercubes, are easily handled in our theory. Given an array of size $2^{m_0} \times 2^{m_1} \times \dots \times 2^{m_d}$, we claim that its elements can be placed at the nodes of a hypercube (of dimension $m_0 + m_1 + \dots + m_d$) such that two “adjacent” data items in the array are placed at neighboring nodes in the hypercube. Here, two data items of the array are *adjacent* if their indices differ in exactly one dimension, and by 1 modulo N , where N is the size of that dimension. (This is called “wrap around” adjacency.)

The following embedding algorithm is described in [34, Section 3.1.2]; it works as follows. If the array has only one dimension with 2^m elements, then we create a gray code sequence, $G\ m$ (see Section 8.4.3). Abbreviate $G\ m$ by g . We place the i^{th} item of the array at the node with label g_i . Adjacent items, at positions i and $i + 1$ (+ is taken modulo $2^m - 1$), are placed at nodes g_i and g_{i+1} which differ in exactly one bit, by the construction.

This idea can be generalized to higher dimensional arrays as follows. Construct gray code sequences for each dimension independently; store the item with index (i_0, i_1, \dots, i_d) at the node $(g_{i_0}; g_{i_1}; \dots; g_{i_d})$ where “;” denotes the concatenations of the bit strings. By definition, adjacent items differ by 1 in exactly one dimension, k . Then, their gray code indices are identical in all dimensions except k and they differ in exactly one bit in dimension k .

We describe a function, em , that embeds an array in a hypercube. Given an array of size $2^{m_0} \times 2^{m_1} \times \dots \times 2^{m_d}$ it permutes its elements to an array $\underbrace{2 \times 2 \times \dots \times 2}_m$, where $m = m_0 + \dots + m_d$, and the permutation preserves array adjacency as described. The algorithm is inspired by the gray code function of Section 8.4.3. In the following, S matches only with a scalar and P with a powerlist.

$$\begin{aligned} em\langle S \rangle &= \langle S \rangle \\ em\langle P \rangle &= em\ P \\ em\langle u \mid v \rangle &= \langle em\ u \rangle \mid \langle em\ (rev\ v) \rangle \end{aligned}$$

The first line is the rule for embedding a single item in 0-dimensional hypercube. The next line, simply, says that an array having length 1 in a dimension can be embedded by ignoring that dimension. The last line says that a non-singleton array can be embedded by embedding the left half of dimension 0 and the reverse of the right half in the two component hypercubes of a larger hypercube.

8.6 Remarks

Related Work

Applying uniform operations on aggregates of data have proved to be extremely powerful in APL [23]; see [3] and [5] for algebras of such operators. One of the earliest attempts at representing data parallel algorithms is in [42]. In their words, “an algorithm... performs a sequence of basic operations on pairs of

data that are successively $2^{(k-1)}$, $2^{(k-2)}$, ..., $2^0 = 1$ locations apart". An algorithm operating on 2^N pieces of data is described as a sequence of N parallel steps of the above form where the k^{th} step, $0 < k \leq N$, applies in parallel a binary operation, OPER, on pairs of data that are $2^{(N-k)}$ apart. They show that this paradigm can be used to describe a large number of known parallel algorithms, and any such algorithm can be efficiently implemented on the Cube Connected Cycle connection structure. Their style of programming was imperative. It is not easy to apply algebraic manipulations to such programs. Their programming paradigm fits in well within our notation. Mou and Hudak[40] and Mou[41] propose a functional notation to describe divide and conquer-type parallel algorithms. Their notation is a vast improvement over Preparata and Vuillemin's in that changing from an imperative style to a functional style of programming allows more succinct expressions and the possibility of algebraic manipulations; the effectiveness of this programming style on a scientific problem may be seen in [53]. They have constructs similar to tie and zip, though they allow unbalanced decompositions of lists. An effective method of programming with vectors has been proposed in [7, 8]. He proposes a small set of "vector-scan" instructions that may be used as primitives in describing parallel algorithms. Unlike our method he is able to control the division of the list and the number of iterations depending on the values of the data items, a necessary ingredient in many scientific problems. Jones and Sheeran[24] have developed a relational algebra for describing circuit components. A circuit component is viewed as a relation and the operators for combining relations are given appropriate interpretations in the circuit domain. Kapur and Subramaniam[25] have implemented the powerlist notation for the purpose of automatic theorem proving. They have proved many of the algorithms in this paper using an inductive theorem prover, called RRL (Rewrite Rule Laboratory), that is based on equality reasoning and rewrite rules. They are now extending their theorem prover so that the similarity constraints on the powerlist constructors do not have to be stated explicitly.

One of the fundamental problems with the powerlist notation is to devise compilation strategies for mapping programs (written in the powerlist notation) to specific architectures. The architecture that is the closest conceptually is the hypercube. Kornerup[31] has developed certain strategies whereby each parallel step in a program is mapped to a constant number of local operations and communications at a hypercube node.

Combinational circuit verification is an area in which the powerlist notation may be fruitfully employed. Adams[1] has proved the correctness of adder circuits using this notation. A ripple-carry adder is typically easy to describe and prove, whereas a carry-lookahead adder is much more difficult. Adams has described both circuits in our notation and proved their equivalence in a remarkably concise fashion. He obtains a succinct description of the carry-lookahead circuit by employing the prefix-sum function (See Section 4.7).

Powerlists of Arbitrary Length

The lengths of the powerlists have been restricted to be of the form 2^n , $n \geq 0$, because we could then develop a simple theory. For handling arbitrary length lists, Steele[48] suggests padding enough “dummy” elements to a list to make its length a power of 2. This scheme has the advantage that we still retain the simple algebraic laws of powerlist. Another approach is based on the observation that any positive integer is either 1 or $2 \times m$ or $2 \times m + 1$, for some positive integer m ; therefore, we deconstruct a non-singleton list of odd length into two lists p, q and an element e , where e is either the first or the middle or the last element. For instance, the following function, *rev*, reverses a list.

$$\begin{aligned} rev \langle x \rangle &= \langle x \rangle \\ rev (p \mid q) &= (rev q) \mid (rev p) \\ rev (p \mid e \mid q) &= (rev q \mid e \mid rev p) \end{aligned}$$

The last line of this definition applies to a non-singleton list of odd length; the list is deconstructed into two lists p, q of equal length and e , the middle element. (We have abused the notation, applying \mid to three arguments). Similarly, the function *lf* for prefix sum may be defined by

$$\begin{aligned} lf \langle x \rangle &= \langle x \rangle \\ lf (p \bowtie q) &= (t^* \oplus p) \bowtie t \\ lf (e \bowtie p \bowtie q) &= e \bowtie (e \oplus (t^* \oplus p)) \bowtie (e \oplus t) \\ &\text{where } t = lf (p \oplus q) \end{aligned}$$

In this definition, the singleton list and lists of even length are treated as before. A list of odd length is deconstructed into e, p, q , where e is the first element of the argument list and $p \bowtie q$ constitutes the remaining portion of the list. For this case, the prefix sum is obtained by appending the element e to the list obtained by applying $e \oplus$ to each element of $lf (p \bowtie q)$; we have used the convention that $(e \oplus L)$ is the list obtained by applying $e \oplus$ to each element of list L .

The Interplay between Sequential and Parallel Computations.

The notation proposed in this paper addresses only a small aspect of parallel computing. Powerlists have proved to be highly successful in expressing computations that are independent of the specific data values; such is the case, for instance, in the Fast Fourier Transform, Batcher merge and prefix sum. Typically, however, parallel and sequential computations are interleaved. While Fast Fourier Transform and Batcher merge represent highly parallel computations, binary search is inherently sequential (there are other parallel search strategies). Gaussian elimination represents a mixture; the computation consists of a sequence of pivoting steps where each step can be applied in parallel. Thus

parallel computations may have to be performed in a certain sequence and the sequence may depend on the data values during a computation. More general methods, as in [7], are then required.

The powerlist notation can be integrated into a language that supports sequential computation. In particular, this notation blends well with ML [38] and LISP[37, 49]. A mixture of linear lists and powerlists can exploit the various combinations of sequential and parallel computing. A powerlist consisting of linear lists as components admits of parallel processing in which each component is processed sequentially. A linear list whose elements are powerlists suggests a sequential computation where each step can be applied in parallel. Powerlists of powerlists allow multidimensional parallel computations, whereas a linear list of linear lists may represent a hierarchy of sequential computations.