Compress-and-Conquer for Optimal Multicore Computing

Zhijing G. Mou
Sinovate, LLC
gmou5813@gmail.com

Hai Liu Paul Hudak
Yale University
{hai.liu,paul.hudak}@yale.edu

Abstract

We propose a programming paradigm called *compress-and-conquer* (CC) that leads to optimal performance on multicore platforms. Given a multicore system of p cores and a problem of size n, the problem is first reduced to p smaller problems, each of which can be solved independently of the others (the *compression* phase). From the solutions to the p problems, a compressed version of the same problem of size $O(p)$ is deduced and solved (the *global* phase). The solution to the original problem is then derived from the solution to the compressed problem together with the solutions of the smaller problems (the *expansion* phase).

The CC paradigm reduces the complexity of multicore programming by allowing the best-known sequential algorithm for a problem to be used in each of the three phases. In this paper we apply the CC paradigm to a range of problems including scan, nested scan, difference equations, banded linear systems, and linear tridiagonal systems. The performance of CC programs is analyzed, and their optimality and linear speedup are proven. Characteristics of the problem space subject to CC are formally examined, and we show that its computational power subsumes that of scan, nested scan, and mapReduce.

The CC paradigm has been implemented in Haskell as a modular, higher-order function, whose constituent functions can be shared by seemingly unrelated problems. This function is compiled into low-level Haskell threads that run on a multicore machine, and performance benchmarks confirm the theoretical analysis.

Categories and Subject Descriptors D.1.3 [*Parallel Programming*]

General Terms Algorithms, Languages, Theory.

Keywords Multicore Programming, Parallel Computing, Programming Paradigm, Functional Programming, Scan, Divide and Conquer, Compress and Conquer

1. Introduction

Parallel programs often introduce certain overheads, such as inter-processor communication, synchronization, and so on. Sometimes these overheads even occur at the algorithmic level. In particular, the total number of operations performed by a parallel algorithm is often greater than that for the best sequential algorithm. We believe that, unlike massively parallel computers, the number of processing units on a multicore system is best considered as a constant, independent of the problem size. It follows that the amount of computation on each core should be of the same order as the complexity of the original problem. Therefore, efficient sequential computation within each core is as crucial as the parallel execution of the program by all cores, in terms of overall performance.

The question we ask, then, is whether we can take advantage of the best-known sequential algorithms in multicore computing. A positive answer to the question might not only lead to efficient multicore computation, but also to a reduction in the complexity of multicore programming, in that a new algorithm does not need to be found if a multicore algorithm can be easily derived from the sequential one.

In this paper, we introduce a new programming paradigm that we call *compress-and-conquer* (CC). Given a multicore system of p cores and a problem of size n, the problem is first reduced to p smaller problems, each of which can be solved independently of the others (the *compression* phase). From the solutions to the p problems, a compressed version of the same problem of size $O(p)$ is deduced and solved (the *global* phase). The solution to the original problem is then derived from the solution to the compressed problem together with the solutions of the smaller problems (the *expansion* phase).

Although this idea sounds simple enough, we have found it fruitful to formalize, analyze, carefully implement, and finally apply the method to a number of non-trivial applications. In particular, our contributions include:

- A description of CC as a high-level algorithmic abstraction (or skeleton) for multicore computing, that demonstrates how a multicore algorithm can be derived from a sequential one.
- A proof of the optimality and linear speedup of CC programs.
- A Haskell library that captures CC as a modular higher-order function. This allows a multicore algorithm to be specified in terms of a small set of constituent functions, many of which can be shared amongst different programs, thus enhancing modularity and promoting code reuse.
- An algorithm for mapping a CC abstraction to a multicore platform, as well as a monadic implementation of this algorithm in Haskell, including the use of mutable arrays.
- Identification of the class of problems subject to the CC abstraction, and an understanding of its limitations, along with a proof that CC subsumes multicore programming models based on scan or mapReduce.
- Application of the CC paradigm to several problems including scan, nested scan, second-order difference equations, banded linear systems, tridiagonal systems, and mapReduce,
- Benchmarks of CC programs for some of the above problems that validate well the theoretical performance results.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
DAMP'10, January 19, 2010, Madrid, Spain.
Copyright © 2010 ACM 978-1-60558-859-9/10/01...$10.00

- A critical comparison of CC to divide-and-conquer, and observations for future work, including the use of a nested form of CC that can be mapped onto hierarchical multicore systems.

The paper is organized as follows. We introduce the notion of CC in Section 2. CC algorithms expressed in Haskell are derived in Section 3 for problems including scan, nested scan, second-order linear difference equations, Fibonacci sequence, banded linear systems, tridiagonal linear systems, and mapReduce. In Section 4 we show how CC programs can be compiled for execution on multicore systems; in particular, how logical data dependencies are mapped to inter-core communications. In Section 5 we give an analysis and proof for the optimality of CC in terms of operation count, communication, and scalability. The benchmarks of some CC programs on multicore systems are also presented. In Section 6 we identify the class of problems subject to the paradigm, and its relation to the computational complexity hierarchy. Some variants of CC are given in Section 7. The relation of CC to divide-and-conquer and related work are discussed in Sections 8 and 9, respectively.

2. The Paradigm

We represent a collection over values of type *a* as an abstract data type *S a*, which can be anything like an array, a list, a tree, a set, etc. Given a function f_s :: *S a* → *S b*, we define the compress-and-conquer (CC) of function f as a higher order function as follows:

DEFINITION 2.1. *The algorithm of compress-and-conquer (CC)*

```
cc :: (∀ a . S a →[S a]) →        – divide
      (∀ a . [S a] →S a) →        – combine
      (S b →S c) →                – compress
      ((S d, S a) →S a) →         – expand
      (S c →S a) →                – pre-communication
      (S b →S d) →                – post-communication
      (S a →S b) →                – sequential function
      S a →S b
cc d c co xp com_g com_h f_s s =
   let seg  = d s
       pre  = map (co . f_s) seg
       core = (d . com_h . f_s . com_g . c) pre
       post = map (f_s . xp) (zip core seg)
   in  c post
```

The computation defined by the CC function can be broken clearly into into three-phases, which we will refer to as *compression*, *global*, and *expansion* phases respectively.

1. Compression phase *map (co .* f_s*) . d* : The input is first divided by d into a number of segments, and function f_s is applied in parallel to each segment, with no inter-dependencies. The results are then compressed by function *co* at each segment. Note that in Def. 2.1 we name the divided segments as *seg*, which is preserved and later retrieved in the expansion phase.
2. Global phase $d\ .\ com_h\ .\ f_s\ .\ com_g\ .\ c$: The compressed segments from the compression phase are first combined by c to become a single collection before passed to the pre-communication function com_g. This is followed by an application of the function f_s, and then a post-communication of com_h. The result is again divided into segments, ready to be distributed back.
3. Expansion phase *c . map (*f_s *. xp) . zip* : The results from the global phase are first zipped with the original input segments, and then expanded by function *xp*. Function f_s is applied again to each segment with no inter-dependencies, and the results are finally combined into one collection.

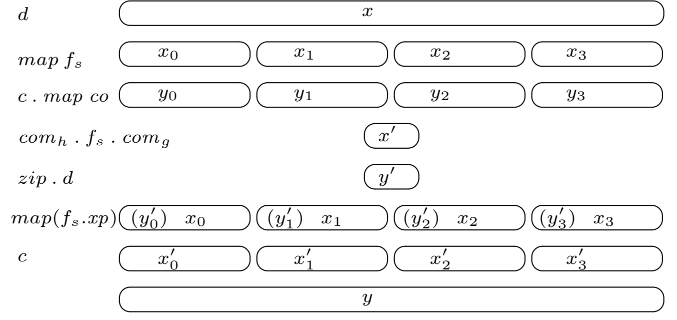

Figure 1. A schematic illustration of a compress-and-conquer algorithm to compute $y = f\ x$ where $f = cc\ d\ c\ co\ xp\ com_g\ com_h\ f_s$ with division arity 4. The first-level oval box represents the input data x, and the last one the output y. The constituent functions to be applied to each level are listed on the left.

A schematic illustration of the CC paradigm is given in Fig. 1. We will refer to the divide, combine, compress, expand, pre- and post-communication, and the sequential function f_s as the *constituents* of compress-and-conquer, They are further explained below:

1. Function *d* :: ∀ *a* . *S a* → [*S a*] divides the given collection into a number of disjoint segments, and the combine function *c* :: ∀ *a* . [*S a*] → *S a* is its left inverse with the property *c* . *d* = *id*. They both are given a polymorphic rank-2 type because we want the division to be independent of the actual values in the collection. For example, list concatenation is polymorphic, whereas the merge in merge-sort and the division in quick-sort are both non-polymorphic.
2. Function *co* :: $S\ b \rightarrow S\ c$ compresses the result after f_s is applied to the input segments before passing them to the global phase. We say that a compress function *co* is *bounded* if there exists a constant k, such that for any s, $|s|/|co\ s| \leq k$, where $|s|$ is the size of collection s. A compress function that is not bounded is *unbounded*. For example, a function that maps any set to a singleton set is an unbounded compress function, which compresses a set of any size to one of size one. In contrast, the compression of a vector that returns all the entries with even indices is bounded, and has a compression ratio of two.
3. The expand function *xp* :: (*S d, S a*) → *S a* takes the results of type *S d* from the global phase, and expands them by modifying the segments from the original input of type *S a*, before passing to the function f_s in the final phase.
4. In the global phase, before f_s is applied to the compressed data, the data is pre-processed by function com_g :: $S\ c \rightarrow S\ a$; then the output from f_s is post-processed by function com_h :: $S\ b \rightarrow S\ d$. These are called the pre- and post-communication functions because they represent the logical data dependency between segments.

We define the following properties of a CC algorithm:

- The *arity* of a CC function is the arity of its divide and combine constituents.
- The *compression ratio* of a CC function is the compression ratio of its compression constituent.
- A CC function has an *unbounded compression ratio* if its compression constituent is unbounded.

- A CC function is *self-similar* if the CC of f_s defines the same function, i.e. $cc\ d\ c\ co\ xp\ com_g\ com_h\ f_s = f_s$, for some *co, xp,* com_g, com_h, and for any *d* and *c*.

As shall be seen in the later sections, functions defined with the above *cc* forms can be mapped to multicore systems and often lead to algorithms with optimal speedups. The CC higher-order form provides a way to specify a multicore algorithm with often very simple constituent functions.

3. Case Studies

In this section we examine the application of CC to a number of common problems. Because these problems all deal with ordered sequences, without loss of generality we use the list type as a concrete representation for *S a*:

type S a = [*a*]

It is important to note that programs written using the list representation are not meant to be efficient implementations, but rather specifications with sufficient detail to guide real implementations over multi-cores that will be discussed in Sec 4.

We also define a few commonly used constituent functions:

```
d :: Int → S a → [S a]
d p l | p == 1      = [l]
      | otherwise = let (m, n) = splitAt (length l `div` p) l
                     in m : d (p - 1) n
c :: Int → [S a] → S a
c p = concat

first, last, last2, bothend :: S a → S a
first     l = take 1 l
first2    l = take 2 l
last      l = drop (length l - 1) l
last2     l = drop (length l - 2) l
bothend  l = first l ++ last l

sr :: a → S a → S a
sr i l = i : take (length l - 1) l
```

Function *d* divides the given sequence into *p* equal-size segments, and *c* is its inverse. Functions *first,first2,last,last2,bothend* are simple constituent functions that extract the first, frst two, last, last two, or both first and last elements from a sequence. Function *sr* shifts the given sequence one position to the right, and fills in the first element with its argument.

3.1 Scan

Scan (or *prefix*) has been considered a powerful parallel and multicore programming construct. Here is a formal definition:

DEFINITION 3.1. *A scan or prefix operation is defined to be a function that maps an input sequence $x_0, x_1, \ldots, x_{n-1}$ with respect to an associative binary operator $\oplus$ to an output of:*

$$x_0, x_0 \oplus x_1, \ldots, x_0 \oplus x_1 \oplus \cdots \oplus x_{n-1}$$

In Haskell, a function called *scanl1* from the *Prelude* already does exactly this computation [16]. So we'll just define our sequential scan as:

scan = *scanl1*

We next show a CC algorithm for scan by providing its simple constituents.

ALGORITHM 3.1. *Scan with respect to an associative binary operator $\oplus$ by compress-and-conquer:*

```
ccScan ⊕ = cc (d p) (c p) last addfirst id (sr 0) (scan ⊕)
  where addfirst ([v], (x : xs)) = v ⊕ x : xs
```

Informally, the *cc* higher-order function takes seven of its constituents, and returns a function that computes the scan with respect to the binary associative operator $\oplus$. It does so by first dividing the input sequence into p segments, and applying the scan over each segment, all segments in parallel. The last elements of the segmented scan are then used to derive a compressed sequence of size p. Scan is then performed over the compressed sequence. The post communication shifts the global result to the right by one position so that the ith result is distributed back to the $(i+1)$th segments, and added to the first element in the original segment by the expand function *addfirst*. A scan is then performed again in parallel to all the segments. All segments are then concatenated to form the final solution (See Figure 2).

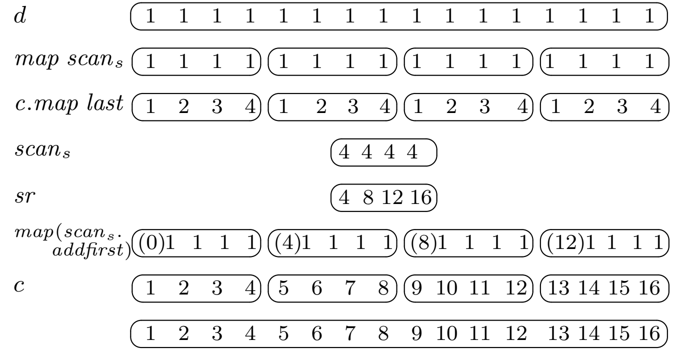

Figure 2. Scan by CC with a sequence of 16 1's, and p = 4.

3.2 Nested Scan

A *nested scan* applies scan to a list of sequences. More formally, we have:

DEFINITION 3.2. *A nested scan with respect to an associate binary operator $\oplus$ is defined in terms of scan (see Def. 3.1):*

nestedScan $\oplus$ = *map* (*scan* $\oplus$)

The solution to nested scan is usually a mapping to the flat scan without involving nested parallelism. We first convert the list of sequences to a flat sequence of pairs with the following function:

```
flat :: [S a] →S (a, Bool)
flat l = zip (concat l) [n == length v | v ← l, n ← [1..length v]]
```

Intuitively, the second component of each pair is a flag indicating whether or not the original element was the last element in the nested sequence. For example:

flat $[[1,2,3],[4,5],[6]] = [(1,\circ),(2,\circ),(3,\bullet),(4,\circ),(5,\bullet),(6,\bullet)]$

where $\bullet$ = *True*, and $\circ$ = *False*. We also define the inverse of *flat* and a lifting function as follows:

```
unflat :: S (a,Bool) → [S a]
unflat l | null l       = []
         | otherwise = let (m, (v : n)) = break snd l
                        in map fst (m ++ [v]) : unflat n

lift :: (a → a → a) → ((a, Bool) → (a, Bool) → (a, Bool))
lift f (x, u) (y, v) = (if u then y else f x y, v)
```

It can be easily verified that if $\oplus$ is associative over type *a*, then *lift* $\oplus$ is associative over type (*a, Bool*).

ALGORITHM 3.2. *Nested scan with respect to an associative binary operator $\oplus$ can be reduced to a flat scan over pairs by*

ccNestedScan $\oplus$ = *unflat* . *ccScan* (*lift* $\oplus$) . *flat*

3.3 Second-Order Linear Difference Equations

In this section, we consider the system of second-order linear difference equations of the following form:

DEFINITION 3.3. *A system of second-order linear difference equations is:*

$$\begin{aligned} & y_0 = c_0 \\ & y_1 = c_1 \\ & y_2 = a_2 y_0 + b_2 y_1 + c_2 \\ & \vdots \\ & y_{n-1} = a_{n-1} y_{n-3} + b_{n-1} y_{n-2} + c_{n-1} \end{aligned} \tag{1}$$

Let us consider a section of (1) corresponding to the variables indexed from s to t, $s < t < n$, denoted by $L[s, t]$:

$$\begin{aligned} y_s \quad & = a_s \quad y_{s-2} + b_s \quad y_{s-1} + c_s \\ y_{s+1} \quad & = a_{s+1} y_{s-1} + b_{s+1} y_s \quad + c_{s+1} \\ y_{s+2} \quad & = a_{s+2} y_s \quad + b_{s+2} y_{s+1} + c_{s+2} \\ & \vdots \\ y_t \quad & = a_t \quad y_{t-2} + b_t \quad y_{t-1} + c_t \end{aligned} \tag{2}$$

In a system of difference equations, we say a variable y_i *depends* on another variable y_j if y_j appears as a term on the right-hand side of its equation; and two variables are *aligned* if they depend on the same variables. We next will align all the variables from y_s to y_t, so that they all depend on the external variables y_{s-2} and y_{s-1}. This can be achieved with the following sequential algorithm:

ALGORITHM 3.3. *A sequential internal solver for a section of second-order difference equation: Given a section $L[s, t]$, where only the first two variables may have external references, and all the other variables refer to variables internal to the section. Let $X = (y_{s-2}, y_{s-1}, 1)$. We define a new sequence of vectors u_i such that $y_i = u_i * X$, where $*$ stands for a point-wise multiplication for vectors.*

$$\begin{aligned} y_s \quad & = u_s * X \\ & = a_s \quad y_{s-2} + b_s \quad y_{s-1} + c_s \\ & = (a_s, b_s, c_s) * X \\ y_{s+1} \quad & = u_{s+1} * X \\ & = a_{s+1} y_{s-1} + b_{s+1} y_s \quad + c_{s+1} \\ & = (a_s b_{s+1}, a_{s+1} + b_s b_{s+1}, c_{s+1} + c_s b_{s+1}) * X \\ \vdots \\ y_t \quad & = u_t * X \\ & = \left(\begin{bmatrix} u_{t-2} \\ u_{t-1} \\ 0\,0\,1 \end{bmatrix} (a_t, b_t, c_t) \right) * X \end{aligned}$$

In Haskell, we write the internal solver as a function mapping from the sequence of (a_i, b_i, c_i) to the sequence of u_i as follows:

```
diff ((a0, b0, c0) : (a1, b1, c1) : xs) = u
  where u0 = (a0, b0, c0)
        u1 = (a0 * b1, a1 + b0 * b1, c1 + c0 * b1)
        u  = u0 : u1 : zipWith3 f u (tail u) xs
        f x y z = (x, y, (0, 0, 1)) ⋉ z
```

where $\ltimes$ is defined to be the operation of multiplying a 3x3 matrix with a vector of size 3.

The above gives a definition of vector sequence u_i, for $s \le i \le t$, and we have successfully aligned all variables from y_s to y_t to the external variables represented by $X = (y_{s-2}, y_{s-1}, 1)$.

Note that *diff* can also be used to solve a complete system of 2nd order linear difference equations where $a_0 = b_0 = a_1 = b_1 = 0$. It doesn't matter how we initialize the two variables in X, *diff* will always return a sequence of $u_i = (0, 0, y_i)$. In this sense, Algo. 3.3 is an algorithm for a generalized form of second-order linear difference equations.

Now consider a system L of n second-order difference equations partitioned into p sections. By applying Algo. 3.3 to each section, we can make all the internal variables of each section align to the last two variables of the previous section. Let L' be a system of equations formed by taking the last two equations from each section, then it is not hard to see, with a little adjustment, what we get is in turn a closed second-order difference equation, with a smaller size of $2p$. We call L' a compressed version of L.

The adjustment needed here is to make the last variable from each section, except the first section, instead of aligning with the last two variables from the previous section, align with the last from the previous, and second last from its own section. This is achieved with the following function:

```
adjustdiff (x : x' : xs) = x : x' : aux xs
  where aux [] = []
        aux ((a, b, c) : (a', b', c') : xs) =
            (a, b, c) : (a", b", c") : aux xs
            where a" = b' - b" * b
                  b" = if a == 0 then 0 else a' / a
                  c" = c' - b" * c
```

Furthermore, solving L' means we have solved the last two variables of each section, therefore the first two variables of the next section can in turn be solved. We'll design an expansion function to properly re-initialize the first two variables in each section, so that they becomes individually solvable by Algo. 3.3.

```
initfirst2 ([(_, _, x), (_, _, x')], (u0 : u1 : xs)) =
  (0, 0, y0) : (0, 0, y1) : xs
  where y0 = (x, x', 1) * u0
        y1 = (x', y0, 1) * u1
```

This lead to the following compress-and-conquer algorithm:

ALGORITHM 3.4. Compress-and-conquer for second-order linear difference equations:

```
ccDiff = cc (d p) (c p) last2 initfirst2 adjustdiff (sr2 (0, 0, 0)) diff
  where sr2 v = sr v . sr v
```

Observe that Fibonacci sequence is no more than a homogeneous second-order difference equations with $c_i = 0$ for $0 \le i \le n$ in Eq. 1, therefore can be solved by CC.

ALGORITHM 3.5. *Since Fibonacci sequence is no more than a special case of second-order linear difference equations, Algo. 3.4 applies.*

3.4 Banded Lower Triangular Linear Systems

DEFINITION 3.4. *A banded lower triangular linear system with bandwidth of two is:*

$$\begin{bmatrix} \dot{a}_0 & & & & \\ \dot{a}_1 & \dot{b}_1 & & & \\ \dot{a}_2 & \dot{b}_2 & \dot{c}_2 & & \\ & \dot{a}_3 & \dot{b}_3 & \dot{c}_3 & \\ & & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} d_0 \\ d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix}$$

By multiplying out the matrix and the vector of unknowns, and some simple algebraic transformation, the above banded linear system becomes a second-order difference equation in the form of (1), where

$$\begin{aligned} &y_0 = d_0/\dot{a} \\ &y_1 = (d_1 - d_0)/\dot{b}_1 \\ &y_2 = -(\dot{b}_2/\dot{c}_2)y_1 - (\dot{a}_2/\dot{c}_2)y_0 + d_2 \\ &\vdots \end{aligned} \tag{3}$$

In other words, a banded linear system is equivalent to a difference equation where the bandwidth equal of the banded system is equal to the order of the difference equations. Algo. 3.4 therefore is also a compress-and-conquer algorithm for banded linear systems of bandwidth two.

ALGORITHM 3.6. *Banded triangular linear systems with bandwidth of two:*

Convert the system to a second-order difference equations by (3), and then apply Algo. 3.4.

In fact, Algo. 3.4 can be easily generalized to linear difference equations of kth order, for arbitrary k, and therefore Algo. 3.6 can also be generalized to solved triangular linear systems with arbitrary bandwidth of k. We choose however to omit the details of the generalization from this paper.

3.5 Tridiagonal Linear Systems

In all the previous case studies, the inter-dependencies between variables are one directional in that if we lay the variables from left to right by their indices, then the dependencies are all from right to left. Tridiagonal linear systems are examples of applications where the dependencies are bi-directional.

The following is a general form for tridiagonal linear system L with n unknowns:

DEFINITION 3.5. *A tridiagonal linear system is:*

$$\begin{bmatrix} b_0 & c_0 & & & & \\ a_1 & b_1 & c_0 & & & \\ & a_2 & b_2 & c_2 & & \\ & & a_3 & b_3 & c_3 & \\ & & & \ddots & \ddots & \ddots \\ & & & & a_{n-1} & b_{n-1} \end{bmatrix} \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_{n-1} \end{bmatrix} = \begin{bmatrix} d_0 \\ d_1 \\ d_2 \\ d_3 \\ \vdots \\ d_{n-1} \end{bmatrix}$$

Note that for a given variable y_i the coefficients a_i and c_i represent its dependency on y_{i-1}, and y_{i+1} respectively in the above standard form. The coefficient a_i and c_i are referred to as the forward and backward dependency coefficients, respectively.

Now let us consider a section $L[s, t]$ of the tridiagonal system consists of the rows corresponds to variables y_s to y_t, where $0 \leq s < t \leq n$.

$$\begin{bmatrix} a_s & b_s & c_s & & & \\ & a_{s+1} & b_{s+1} & c_{s+1} & & \\ & & \ddots & \ddots & \ddots & \\ & & & a_t & b_t & c_t \end{bmatrix} \begin{bmatrix} y_s \\ y_{s+1} \\ \vdots \\ y_t \end{bmatrix} = \begin{bmatrix} d_s \\ d_{s+1} \\ \vdots \\ d_t \end{bmatrix}$$

A variable y_i is said to be *forward (backward) aligned* with y_j if they are forward (backward) dependent on the same variables. They are said to be *aligned* if they are both forward and backward aligned. Hence, no two variables are aligned in the above diagram.

Variable can be aligned by Gaussian elimination. For example, The variable y_{s+1} can be forward aligned with y_s by multiply the row for y_s by $-a_{s+1}/b_s$, and then add it to the row for y_{s+1}. We repeat this process for every row except the row for y_s in $L[s, t]$, which leads to the following diagram:

$$\begin{bmatrix} a'_s & b'_s & c'_s & & \\ a'_{s+1} & & b'_{s+1} & c'_{s+1} & \\ \vdots & & & \ddots & \ddots \\ a'_t & & & b'_t & c'_t \end{bmatrix} \begin{bmatrix} y_s \\ y_{s+1} \\ \vdots \\ y_t \end{bmatrix} = \begin{bmatrix} d'_s \\ d'_{s+1} \\ \vdots \\ d'_t \end{bmatrix}$$

Now variables y_{s+1} to y_t are forward aligned with y_s, which means that the coefficients of $a'_s, a'_{s+1}, \ldots, a'_t$ are in the same column in the matrix. We can write the forward alignment as a function that takes a sequence of (a_i, b_i, c_i, d_i) and returns the modified co-efficients (a'_i, b'_i, c'_i, d'_i) as follows:

```
forward [] = []
forward (x : xs) = u
  where u = norm x : zipWith f xs u
        f (a, b, c, d) (a', b', c', d') =
          norm (-a' * a, b - c' * a, c, d - d' * a)
norm (a, b, c, d) = (a / b, 1, c / b, d / b)
```

Note that in the process we also normalize every row so that the co-efficients on the diagonal of the matrix (all the b_i) become 1.

With the forward alignment in place, we can use a similar process to backward align variable y_{t-1} with y_t, and so on, which leads to the following diagram:

$$\begin{bmatrix} a''_s & b''_s & & & & c''_s \\ a''_{s+1} & & b''_{s+1} & & & c''_{s+1} \\ \vdots & & & \ddots & & \vdots \\ a''_{t-1} & & & & b''_{t-1} & c''_{t-1} \\ a''_t & & & & b''_t & c''_t \end{bmatrix} \begin{bmatrix} y_s \\ y_{s+1} \\ \vdots \\ y_{t-1} \\ y_t \end{bmatrix} = \begin{bmatrix} d''_s \\ d''_{s+1} \\ \vdots \\ d''_{t-1} \\ d''_t \end{bmatrix}$$

Note that all variables y_s to y_t are now both forward and backward aligned. We can write the backward alignment function in a similar manner as follows:

```
backward [] = []
backward u = reverse v
  where (x : xs) = reverse u
        v = x : zipWith f xs v
        f (a, b, c, d) (a', b', c', d') =
          (a - a' * c, b, - c' * c, d - d' * c)
```

We consider a tridiagonal system is solved if only diagonal coefficients are left in the matrix. Obviously, if $a_s = c_t = 0$, the system $L[s, t]$ is completely solved after forward and backward alignments. When a_s or c_t are not zeros, however, we shall only align the inner block $L[s + 1, t - 1]$, and adjust the boundary rows for y_s and y_t to align inward like this:

```
adjust l = let [(a0, b0, c0, d0), (a1, b1, c1, d1)] = first2 l
               [(a2, b2, c2, d2), (a3, b3, c3, d3)] = last2 l
           in [(a0, b0 − a1 ∗ c0, −c1 ∗ c0, d0 − d1 ∗ c0)] ++
              middle l ++
              [(−a2 ∗ a3, b3 − c2 ∗ a3, c3, d3 − d2 ∗ a3)]
```

As as result from this adjustment, we effectively obtain a diagram of the following shape for $L[s, t]$ when $a_s \neq 0$ or $c_t \neq 0$:

$$\begin{bmatrix} a''_s & b''_s & & & c''_s & \\ & a''_{s+1} & b''_{t+1} & & c''_{s+1} & \\ & \vdots & & \ddots & \vdots & \\ & a''_{t-1} & & b''_{t-1} & c''_{t-1} & \\ & a''_t & & & b''_t & c''_t \end{bmatrix} \begin{bmatrix} y_s \\ y_{s+1} \\ \vdots \\ y_{t-1} \\ y_t \end{bmatrix} = \begin{bmatrix} d''_s \\ d''_{s+1} \\ \vdots \\ d''_{t-1} \\ d''_t \end{bmatrix}$$

ALGORITHM 3.7. *A sequential internal solver for a section of tridiagonal linear systems is a composition of the forward and backward alignment, and the adjustment function:*

$$\begin{aligned}
&\textit{trid}\ [\,] = [\,]\\
&\textit{trid}\ l = \text{case}\ (a, c)\ \text{of}\\
&\quad (0,0) \rightarrow \textit{backward}\ (\textit{forward}\ l)\\
&\quad _ \qquad \rightarrow \textit{adjust}\ ([x] +\!\!+\ \textit{backward}\ (\textit{forward}\ (\textit{middle}\ l)) +\!\!+\ [y])\\
&\quad \text{where}\ [x@(a,_,_,_), y@(_,_,c,_)] = \textit{bothend}\ l
\end{aligned}$$

Now if we divide a tridiagonal system into p sections, and apply Algo. 3.7 to each section, they are then all internally solved. In Figure 3, we show the non-zero coefficients in the matrix after internally solving all sections for an example case where $n = 16$ and $p = 4$.

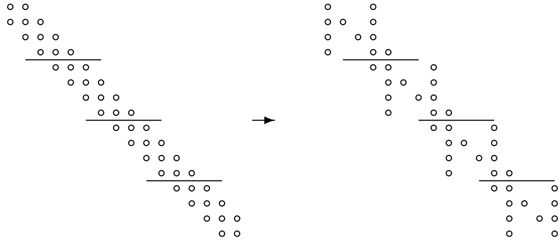

Figure 3. A tridiagonal linear system of size n (=16) divided into p (=4) sections, each section internally solved.

By focusing on the first and last variables in all sections after the internal solver *trid* is applied, one realizes that they in turn form a compressed tridiagonal system of size $2p$. This compressed system can in turn be solved by *trid*. The solution of the compressed tridiagonal system can be plugged back to each section, and each section can then be completely solved independently. This leads to the following compress-and-conquer algorithm for tridiagonal linear systems:

ALGORITHM 3.8. Compress-and-conquer algorithm for tridiagonal linear systems

$$\begin{aligned}
&\textit{ccTrid} = \textit{cc}\ (d\ p)\ (c\ p)\ \textit{bothend replace id id trid}\\
&\quad \text{where}\ \textit{replace}\ ([x, y], l) = [x] +\!\!+\ \textit{middle}\ l +\!\!+\ [y]
\end{aligned}$$

3.6 MapReduce

DEFINITION 3.6. *MapReduce is the functional composition of the map and reduce:*

$$\textit{mapReduce}\ f \oplus = \textit{reduce} \oplus .\ \textit{map}\ f$$

where reduce with respect to an associative binary operator $\oplus$ is a function that maps a non-empty sequence $x_0, x_1, \ldots, x_{n-1}$ to a single value of $x_0 \oplus x_1 \oplus \cdots \oplus x_{n-1}$.

J. Dean and S. Ghemawat introduced *mapReduce* in [11] as a separate programming construct, gave distributed implementations, and showed it applies to many search engine problems.

The problem of *mapReduce* can be said to be an inherently simpler problem than any of the problems we have considered so far and can be computed by a compress-and-conquer where the post-phase is not needed. We therefore introduce a new and simpler version of compress-and-conquer, which we call *pre-CC*, for it contains only the pre-phase of the more general CC form as in Def. 2.1:

$$cc_{pre}\ d\ c\ co\ f = co\ .\ c\ .\ map\ (co\ .\ f)\ .\ d$$

We then have the following simple algorithm for the parallel version of mapReduce:

ALGORITHM 3.9. *MapReduce with respect to an associative binary operator $\oplus$ and a function f is defined in terms of pre-CC:*

$$\textit{ccMapReduce}\ f \oplus = cc_{pre}\ (d\ p)\ (c\ p)\ (\textit{reduce} \oplus)\ (\textit{map}\ f)$$

4. Implementation

4.1 Operational Mapping

Implementation of compress-and-conquer algorithms on multicore systems is fairly straightforward. The work done in compression and expansion phases can be easily mapped onto p threads or processors in parallel. We can certainly use a more compact representation than lists, but more fundamentally, the specification of CC as given in Def. 2.1 is inefficient on today's dominant CPU architectures due to the immutability implied by referential transparency, which prevents destructive updates. Also, the divide and combine functions should just share the original input data instead of making new copies of them, and the order and the arity of the divide function needs to be consistent with combine.

For the above reasons, we move to a monadic form of compress-and-conquer in Haskell [16]:

DEFINITION 4.1. *The implementation of monadic compress-and-conquer (cc_m):*

$$\begin{aligned}
&cc_m :: \textit{Monad}\ m \Rightarrow\\
&\qquad (\forall a.\ ([S\ a] \rightarrow m())\\
&\qquad\qquad \rightarrow S\ a \rightarrow m(S\ a)) \rightarrow && \text{– divide then combine}\\
&\qquad (\forall a.\ (S\ a \rightarrow m())\\
&\qquad\qquad \rightarrow [S\ a] \rightarrow m[S\ a]) \rightarrow && \text{– combine then divide}\\
&\qquad (S\ a \rightarrow m(S\ b)) \rightarrow && \text{– compress}\\
&\qquad ((S\ c, S\ a) \rightarrow m(S\ a)) \rightarrow && \text{– expand}\\
&\qquad (S\ b \rightarrow m(S\ a)) \rightarrow && \text{– pre-core}\\
&\qquad (S\ a \rightarrow m(S\ c)) \rightarrow && \text{– post-core}\\
&\qquad (S\ a \rightarrow m(S\ a)) \rightarrow && \text{– sequential}\\
&\qquad S\ a \rightarrow m(S\ a)\\
&cc_m\ dc\ cd\ co\ xp\ g\ h\ f_s = dc\ aux\\
&\quad \text{where}\ aux\ seg = \text{do}\\
&\qquad\quad pre\ \leftarrow parmap\ (co \cdot f_s \cdot dup)\ seg\\
&\qquad\quad core \leftarrow cd\ (h \cdot f_s \cdot g)\ pre\\
&\qquad\quad parmap\ (f_s \cdot xp)\ (zip\ core\ seg)\\
&\qquad (f \cdot g)\ x = g\ x \ggg\!= f
\end{aligned}$$

In this program, we intentionally define a composition operator $(\cdot)$ to be the monadic counterpart of function composition so that our implementation of cc_m closely matches the specification of CC in Def. 2.1. Further explanations are given below:

1. In order to do destructive updates, we must now make the sequential function f_s return the same collection type as its input. This affects the overall types of cc_m and its constituent functions.
2. We pair up the divide and combine functions as either a single divide-then-combine or combine-then-divide operation. Both are now higher-order functions that take as argument a function that can update the original data in place, but can not change the structure of them.
3. Because the original CC algorithm requires the input collection to remain unchanged until the expansion phase, we must use $dup :: S\ a \rightarrow m(S\ a)$ to create a local copy of the segment during the compression phase.
4. The original *map* function is changed to a monadic *parmap* $:: (a \rightarrow m\ b) \rightarrow [a] \rightarrow m[b]$ that spawns off a system thread for each segment, and only returns when all threads are done.

In our actual implementation, we choose to define the concrete collection type as an unboxed mutable array as follows in order to minimize computation overhead:

data S a = Arr (IOUArray Int a)	– shared array
Int	– lower bound
Int	– upper bound

This definition leads to straightforward implementation of both *cd* and *dc* by sharing the original array without duplicating them. All the constituent functions used in the specification of our algorithms must also be modified to operate on arrays, with direct indices and destructive updates. We omit such details here.

Similarly, the sequential algorithms for all the applications we considered in Section 3 need to be modified to a monadic version that works on the concrete $S\ a$ type defined above, while their compress-and-conquer algorithms require little change except moving from cc to cc_m.

Among the different parallel facilities that GHC (Glasgow Haskell Compiler) provides, the lightweight thread library becomes a natural choice because the IO monad implements destructive update. So in other words, we choose the monad m in Def. 4.1 to be just *IO*, and *parmap* is implemented using *forkIO*. We also choose the division parameter p to match the number of cores in the hardware so that the original array is split into p segments, and consequently *parmap* spawns exactly p system threads. We rely on the operating system to balance system threads among multiple cores.

4.2 Inter-Core Communications

With the mapping of CC algorithms to multicore systems given in Section 4.1, and if we assume the divided segments reside locally to each processor, we can see that there are two, and only two, constituent functions in a CC algorithm that involve inter-core communications: the results from co at the end the compression phase are moved over to the global phase, and after the global phase, the results are moved back to each processor as input to the expand function. The remaining constituent functions are mapped to local operations. Note that the constituent functions com_g and com_h are referred to as communication functions, not because they are to be mapped into inter-core communications at the implementation level, but rather they realize the dependency relations between different sections in the logic domain.

Let $S = (P_0, \ldots, P_{p-1})$ be a multicore system with p cores used by a CC algorithm, and, without loss of generality, P_0 be the appointed core for the global computation, then by the mapping of $parmap$ from Section 4.1, one can see that

- At the end of the compression phase, each P_i ,for $0 < i < p$, sends one piece of data to P_0.
- At the beginning of the expansion phase, each P_i receives a piece of data from P_0.

If we go beyond a Haskell implementation, in the *Message Passing Interface* (MPI) [7], there are two supported communication patterns, *gather* and *scatter*, that perform precisely the above two operations respectively. It is therefore straightforward to support the communication in CC algorithms with MPI. Other options, including MP, PThreads, Intel's Thread Building Blocks [10], and Microsoft's Parallel Task Library [13], can all be used as well.

5. Performance Analysis

Since parallel programs generally incur some overhead over the best known sequential counterparts for the same problems, it is a good practice to understand and quantify the overhead asymptotically. In this section, we show that the overhead of CC algorithms in both operation and communication aspects are minimum, which also translates to linear speedups on multicore systems.

5.1 Operation Optimality

Given a program P, its *operation complexity*, written $\psi(P)$, is the total number of operations that P performs, as a function of the problem size. We say two programs P_1 and P_2 are *consistent* with each other, written $P_1 \sim P_2$, in operation complexity if and only if $\psi(P_1) = \Theta(\psi(P_2))$ [1].

THEOREM 5.1. *Let f be a CC program with base function f_s (see Def. 2.1), then $f \sim f_s$. In other words, a CC program is consistent with its base function.*
Proof: besides the sequential base function, all other constituents in the CC program takes time independent of problem size.

Given a problem f, its *operation complexity*, written $\phi(f)$, is the minimum number of operations f inherently requires, as a function of the problem size. We say a program F that solves problem f is *operation optimal* for f, written $F \propto_o f$ if and only if $\psi(F) = O(\phi(f))$.

It follows from the above definition and Theorem 5.1 that:

THEOREM 5.2. *Given a problem f, a CC program F that solves f, and the sequential base function f_s for F, then $F \propto_o f$ if and only if $f_s \propto_o f$.*

The above theorem gives a convenient way to test for the operation optimality of CC programs. From which one can easily verify that:

THEOREM 5.3. *The CC algorithms Algo. 3.1 for scan, Algo. 3.2 for nested scan, Algo. 3.4 for second-order difference equations, Algo. 3.5 for Fibonacci sequence, Algo. 3.6 for banded linear systems, and Algo. 3.8 for tridiagonal linear systems, are operation optimal.*

5.2 Communication Optimality

Given a multicore program P, its *communication complexity*, written $\delta(P)$, is the total number of inter-core communications that P performs, as a function of the number of cores p. We say two programs P_1 and P_2 are *consistent in communication*, written $P_1 \approx P_2$, in communication complexity if and only if $\delta(P_1) = \Theta(\delta(P_2))$.

Given a problem f over input X partitioned into p disjoint and non-empty subsets of X, we say its *communication complexity*, written $\gamma(f)$, is the minimum number of references crossing the partitions that f inherently requires, as a function of the number of partitions m. We say a multicore program F solving problem f is *communication optimal* for f, written F $\propto_c f$ if and only if $\delta(F) = O(\gamma(f))$.

THEOREM 5.4. *The CC algorithms Algo. 3.1 for scan, Algo. 3.2 for nested scan, Algo. 3.4 for second-order difference equations, Algo. 3.5 for Fibonacci sequence, Algo. 3.6 for banded linear systems, Algo, 3.8 for tridiagonal linear systems are all communication optimal.*
Proof: Let f be any of the above problems, X the input for f. Suppose X is partitioned into any p disjoint and non-empty blocks. Since the final solution of f over X depends on at least one piece of data in each of the m blocks, the communication complexity of f, $\gamma(f(p)) = \Omega(p)$ [2]. But the CC algorithm for f has communication complexity $\delta f(p) = O(p)$. Therefore, the CC algorithm for f is communication optimal.

[1] $f = \Theta(g)$ if and only if $f = O(g)$ and $g = O(f)$

[2] given f and g, f is said to be at least of the order of g, written $f = \Omega(g)$, if $g = O(f)$

5.3 Linear Speedups

Let f be a program, $T(f, n, p)$ the time to carry out f on input size n and p cores. Then the *speedup* of f is:

$$S(n) = T(f, n, 1)/T(f, n, p) \quad (4)$$

It follows that:

THEOREM 5.5. *Let p be the number of cores, n size of the input. If $p = o(n)$ [3], then, the CC algorithms Algo. 3.1 for scan, Algo. 3.2 for nested scan, Algo. 3.4 for second-order difference equations, Algo. 3.5 for Fibonacci sequence, Algo. 3.6 for banded linear systems, Algo, 3.8 for tridiagonal linear systems have asymptotical speed up linear to the number of cores p.*
Proof: In all the above algorithms, the compression and expansion phases take $O(n/p)$ time, and the global phase takes $O(p)$ time. The total time is then $T(f, n, p) = O(n/p) + O(p)$. Since $p = o(n), T(f, n, p) = O(n/p)$. By (4):

$$S(n) = T(f, n, 1)/T(f, n, p) = O(n)/O(n/p) = O(p) \quad (5)$$

Also to be observed is:

THEOREM 5.6. *The computational time of the global phase in a CC algorithm is a function of the number of cores p, and independent of the size n of the problem.*

The above implies that if the sequential base constituent of a CC algorithm is an optimal one sequentially, then the CC algorithm is also an optimal multicore program in the sense that, (1) it is a consistent algorithm, and (2) it has linear speedup.

In Figure 4, we plot the speed-up curve of some CC programs in Haskell on an Intel Xeon machine runing Linux OS with seven cores available to us. We omit the benchmark for nested scan because it is being implemented in terms of a single flat scan. All programs run entirely in memory, and we measure the speed by calculating the wall time each program takes from start to finish. Observe that the speedups for the three different problems are all nearly perfectly linear to the number of cores used for the computation.

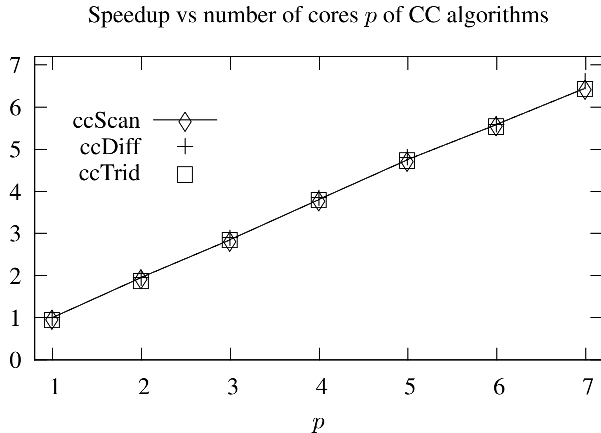

Figure 4. Speedup curve of CC programs in Haskell for scan, second-order difference equation, and tridiagonal problem for $n = 10^6$ on multicore system with seven cores.

Theorem 5.5 may appear to be a direct violation to Amdahl's Law [1] or Gustafson's Law [8]. There is a simple explanation for this. Both laws assume some fixed percentage of either the parallel or sequential portion of a program. Since this is not a valid assumption for the problems we consider in this paper, neither Amdahl's nor Gustafson's Law is relevant here.

[3] $f(x) = o(g(x))$ if $\lim_{x\to\inf} f(x)/g(x) = 0$

6. Characteristics

In the above section, it is shown that, for a broad range of problems, CC paradigm can deliver multicore solutions optimal in computation and communication with linear speed up. It is however unclear what the common characteristics are of the problems that are subject to the CC programming paradigm.

To answer the above question, let us first introduce the notion of the $\mathcal{CC}$ *class*.

DEFINITION 6.1. *A problem is in the class $\mathcal{CC}$ if and only it is subject to the CC form of (Def. 2.1) with an unbounded compression ratio (Section 2).*

It follows that:

THEOREM 6.1. *Scan, nested scan, second-order linear difference equations, Fibonacci sequence, banded linear triangular system of bandwidth two, tridiagonal linear systems are in the class $\mathcal{CC}$.*

To characterize problems in $\mathcal{CC}$, we need the following notions:

DEFINITION 6.2. *Let F be a function over input X. The reference graph of F is the pair $G = (V, R)$, where $V = \{x \mid x \in X\}$, and R is the binary relation, where $x_1\ R\ x_2$ if and only if x_1 refers to x_2 in F.*

For instance, the reference graph for the problem of second-order difference equations is a chain of vertices, each of which, with the exception of the first two, has two directed edges connecting it to the two previous ones respectively. Since this binary relation generally is not symmetric, the graph is directed.

Given a graph $G = (V, R)$, a *cut* is a binary partition of the vertices, and the *size of a cut* is the number of edges between the two partitions. A cut is *maximum* if its size is larger than any other cut.

Now we are in a position to identify a necessary condition for problems to be in the class $\mathcal{CC}$:

THEOREM 6.2. *Let f be a problem in $\mathcal{CC}$, then there exist a reference graph $G = (V, R)$ for f with maximum cut independent of $|V|$, where $|V|$ denotes the cardinality of V.*
Proof: suppose this is not the case, we can then use the reference graph as defined by the CC algorithm for that of f. This graph however has maximum cut bounded by a constant, leading to a contradiction.

The problem of second-order difference equations, for instance, has a reference graph that meets the above condition.

It should come as no surprise that all problems are not known to possess reference graphs with constant bounded maximum cut as required by Theorem 6.2. FFT and Bitonic Sort are examples of such problems.

Next, we show that the class $\mathcal{CC}$ is characterized not only by the property of the reference graphs, but also by the complexity classes:

THEOREM 6.3. *Let $\mathcal{L}$ be the class of problems with computational complexity of $O(n)$ [4], where n is the size of the problem. Then $\mathcal{CC} \subset \mathcal{L}$.*
Proof: suppose there is a problem $f \in \mathcal{CC}$, and $f \notin \mathcal{L}$.

[4] Here, the $O(n)$ refers to the linear complexity of a problem on a Turing machine.

Let $T(f, n, 1) = O(g(n))$, where g is not linear to n, f_s the base function of f. The time to compute f_s on each core will be $g(n/p)$. If we simulate the CC program for p cores on one core, the total time will be $O(mg(n/m))$. Since g is more than linear with n, it follows that $O(mg(n/m)) < O(g(n))$, which leads to a contradiction.

Theorem 6.2 and 6.3 point out rather severe limitations on the power of the compress-and-conquer paradigm. However, there are problems that, though not in themselves in the class $\mathcal{CC}$, contain component(s) which are. Matrix multiplication, for instance, is clearly not in the class $\mathcal{L}$, however, its main component, the inner product of a row and a column from the two factor matrices, is in the class $\mathcal{CC}$ and can indeed be computed with a CC program.

7. Variations and Generalizations

7.1 Parallelized Core-Phase

Observe that in the CC form of Def. 2.1 we have chosen to apply the sequential base function over the compressed problem during the global phase, and as a result, the global phase computation is mapped into the internal computation inside a single appointed core (P_0, see Section 4). Alternatively, one could choose a parallel program for the global phase. It can be shown, however, unless the number of cores is sufficiently large, the alternative parallel approach brings no benefit to performance, but only complicates the programming requirement. For if one goes that way, he must provide a separate parallel version of the base function in addition to the sequential version which is shared in all the three phases under the proposed scheme.

7.2 Specialized Sequential Function

An interesting aspect of CC is that the sequential function f_s is applied three times, one in each of the three phases:

1. In compression phase, f_s only partially solves each segment of the original input data;
2. The compressed results form a much smaller problem in the global phase, which is completely solved by f_s;
3. The solution to the compressed problem is expanded to modify each segment of the original data, which then are completely solved by f_s.

For this reason, we'll call f_s the *generalized solver* for a given problem. But in order to re-use the same f_s, we have to retain the original data until the last phase. A consequence made more apparent by the monadic cc_m is that in the compression phase it has to make copies of the input segments otherwise f_s would modify them in place. This is of course an implementation issue that can be addressed, for instance, by some fusion technique. A more fundamental question is: *can we re-use the result of f_s from the compression phase without having to keep the original data around?*

The answer is *yes*. Instead of relying on just one f_s for all phases, we can take another sequential function g_s that we call a *specialized solver*, and formulate a different CC algorithm below:

cc' d c co xp com_g com_h f_s g_s = *post* . *first core* . *pre*
  where *pre* = *unzip* . *map* ((*co* × *id*) . f_s) . *d*
    core = *d* . com_h . f_s . com_g . *c*
    post = *c* . *map* (g_s . *xp*) . (*uncurry zip*)
    first f (*x*, *y*) = (*f x*, *y*)
    f × *g x* = (*f x*, *g x*)

Just like the original *cc* in Def. 2.1, *cc'* still contains three phases, but in the compression phase, it actually passes the results from function f_s directly to the expansion phase, and function g_s would pick up from where f_s has left and work out a complete solution with the expanded information obtained from the global phase.

In terms of complexity, cc' is on the same order as cc. But in an actual implementation, it may perform better because the specialized solver g_s may require less computation steps than the generalized solver f_s since it already has a partial solution to start with. Theoretically, however, we still prefer the original CC formulation in Def. 2.1 which is easier to reason about for its simplicity.

7.3 Higher-Order CC

A compress-and-conquer with a sequential base function is said to be of *first order*. Inductively, a compress-and-conquer is said to be of a $(k+1)$-th order CC algorithm if its base function f_s is a k*-th order compress-and-conquer*.

Let us consider a second-order compress-and-conquer, with arity n at top level, m the bottom level. It can be mapped to a multicore system with n interconnected nodes, each with m cores. It is easy to show that

THEOREM 7.1.

(1) A second $(k+1)$-th order CC is operation and communication optimal if and only if its (k-th order) base function is.

(2) The speedup of a second-order compress-and-conquer with arities n and m at top and base levels respectively mapped to n nodes with m cores is respectively linear to n and m.

Observe that second-order CC form provides a simple and elegant framework to program hierarchical systems with multiple nodes of multicore units with guaranteed optimal performance.

It should also be obvious the above theorem can be generalized to CC algorithms with orders greater than two.

8. Relation to Divide-and-Conquer

Divide-and-conquer (DC) has been shown to be one of the most effective paradigms for deriving elegant and efficient parallel solutions to a wide variety of problems [6, 14, 15]. Both DC and CC solves a problem by dividing it into sub-problems. However, the arity of the division in DC is usually some small constant such as two, while CC uses division with arity variable in the number of processing units; DC is recursive, while CC is not; and a DC algorithms usually is a different algorithm from the best-known sequential counterpart altogether, while a CC algorithm is always derived from a sequential algorithm for the same problem.

It should also be pointed out that the two paradigms are not equivalent in their computational power. Given Theorem 6.3, the computational power of compress-and-conquer is strictly weaker than that of divide-and-conquer.

Also note that the two paradigms do not necessarily lead to the same performance. Scan, for example, although the problem has $O(n)$ operation complexity, a DC algorithm would require $O(nlog(n))$ operations [15]. In contrast, as shown in Theorem 5.1, a CC algorithm would be operation optimal.

Finally, we would like to point out that automatic transformation between DC and CC programs is possible under certain conditions. Our previous work on divide-and-conquer introduced the notion of *pre-* and *post- morphism* as algebraic models for DC, and it was pointed out that a broad range of scientific problems can be solved with three types of communications, namely, last-k, correspondent, and mirror-image [14, 15]. It can be shown that a postmorphism [14, 15] algorithm with last-k communication can be automatically transformed into a CC program, and vice versa, which limited by space must be elaborated elsewhere.

9. Related Work

Much effort has been made to support high-level programming for multicore computing. Some noticeable examples are the Threading Building Blocks from Intel [18], Parallel Task Library from Microsoft, and the Data Parallel Haskell project [5] from the functional programming community. The CC paradigm proposed here differs from any of the above approaches in a number of ways. Firstly, it does not expose any of the mechanisms related to multicore architecture such as thread, mutex, and task queues. Secondly, its does not expose any imperative constructs such parallel-for or parallel loops. Finally, instead of relying on programming constructs such as reduce and scan, it provides a more general form from which the constructs can be derived.

Solving a problem through compression is not an entirely new idea. There is a known technique in parallel computing, referred to as *odd-even reduction*. Ladner and Fischer [12] used this technique in an elegant parallel scan algorithm. With odd-even reduction, a problem is recursively reduced in size by a factor of two. As a result, the number of steps required is logarithmic to the size of the problem during both the reduction and expansion phase. In contrast, the CC paradigm has unbounded compression ratio, and takes one step during both the compression and expansion. Another obvious difference is that an algorithm based on odd-even reduction is totally a different algorithm from its sequential counterpart, while a CC algorithm employs the sequential counterpart as the core of its computation.

Nested data parallelism has been shown to be an expressive and effective approach to multicore programming [9, 17], which can be traced back to work on the language NESL and nested scan [3, 4]. From a data structure point of view, both Data Parallel Haskell and compress-and-conquer introduce new kinds of array operations. The two approaches however have salient differences in nature. First of all, the division of arrays in the former are non-polymorphic in that the result depends on the values of the array entries through the use of array comprehension (e.g. the division used in quicksort), while in the latter, polymorphic structural operations are of fundamentally importance to the paradigm (non-polymorphic operations can be implemented with polymorphic operations). Secondly, in spite of a large number of primitives built into the parallel arrays of the former, data communication is completely hidden, and programmers have to trust the compiler to do a good job of balancing tasks. In contrast, communication in the latter is a first-class citizen. Thirdly, monadic composition (in its comprehension form) is the main theme in the former, while higher-order functional forms are the center pieces of the latter.

M. Cole et. al. with their work on programming skeletons have shown how higher-order functions can be adapted to work with non-declarative languages for the purpose of parallel programming [2, 6]. Under their framework, as a higher-order form, compress-and-conquer can be considered as another algorithmic skeleton, different but related to divide-and-conquer, which they have identified as an important parallel algorithmic skeleton.

10. Conclusion

We have proposed CC as an efficient paradigm for multicore computation, and showed how it can be implemented using Haskell higher-order functions. The expressive power of the paradigm was illustrated with its application to a number of problems including scan, nested scan, difference equations, banded linear systems, and linear tridiagonal systems. The optimality of CC programs was proven and confirmed by the benchmarks of the CC programs on multicore machine. Besides the linear speedup, the CC paradigm reduces the complexity of multlicore programming by allowing a sequential program to be used as the core component of the multicore program. While not all problems are subject to the paradgim, the computational power is shown to subsume that of scan, nested scan, and mapReduce.

Acknowledgement This research was supported in part by a grant from Microsoft Research, and by NSF grant CCF-0811665.

References

[1] G. Amdahl. Validity of the single-processor approach to achieving large-scale computing capabilities. *Proceedings of AFIPS Conference*, pages 483–485, 1967.

[2] S. G. Anne Benoit, Murray Cole and J. Hillston. Why skeletal parallel programming matters. In *Proceedings of Euro-Par 2004*, page 37, 2004.

[3] G. Blelloch. Scans as primitive parallel operations. In *International Conference on Parallel Processing*, 1987.

[4] G. Blelloch. Programming parallel algorithms. *Communication of the ACM*, 39(3), March 1996.

[5] M. M. Chakravarty, R. Leshchinsky, S. Peyton Jones, G. Keller, and S. Marlow. Data parallel Haskell. In *DAMP'07*, November 2007.

[6] M. Cole. Algorithmic skeletons: Structured management of parallel computation. 1989.

[7] W. Gropp and et al. Mpich2 user's guide. *Mathematics and Computer Science Division, Argonne national lab*, November 2004.

[8] J. Gustafson. Reevaluating Amdahl's law. *Communication of the ACM*, 21(5):532–533, 1988.

[9] T. Harris and S. Singh. Feedback directed implicit parallelism. In *International Conference on Functional Programming*, Oct 2007.

[10] Intel. Intel 64 and IA-32 architectures software developer's manual, August 2007.

[11] J.Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. 2004.

[12] R. E. Ladner and M. J. Fischer. Parallel prefix computation. *Journal of the ACM*, 27(4):831–838, 1980.

[13] D. Leijen and J. Hall. Optimize managed code for multi-core machines. *MSDN Magazine*, October 2007.

[14] Z. G. Mou. *A Formal Model for Divide-and-Conquer and Its Parallel Realization*. PhD thesis, Yale University, May 1990.

[15] Z. G. Mou and P. Hudak. An algebraic model for divide-and-conquer algorithms and its parallelism. *The Journal of Supercomputing*, 2(3): 257–278, November 1988.

[16] S. Peyton Jones, editor. *Haskell 98 Language and Libraries – The Revised Report*. Cambridge University Press, Cambridge, England, 2003.

[17] S. Peyton Jones and et. al. Harnessing the multicores: Nested data parallelism in Haskell. In *Foundations of Software and Theoretical Computer Science*, Bangalore 2008.

[18] J. Reinders. *Intel Threading Building Blocks*. O'Reilly, 2007.