ViC*: A Compiler for Virtual-Memory C*
(Extended Abstract)

Alex Colvin [*]
Dartmouth College
Department of Computer Science

Thomas H. Cormen [†]
Dartmouth College
Department of Computer Science

Abstract

This paper describes the functionality of ViC, a compiler for a variant of the data-parallel language C* with support for out-of-core data. The compiler translates C* programs with shapes declared* outofcore, *which describe parallel data stored on disk. The compiler output is a SPMD-style program in standard C with I/O and library calls added to efficiently access out-of-core parallel data. The ViC* compiler also applies several program transformations to improve out-of-core data access.*

Keywords: compilers, parallel I/O, C*, parallel disk systems, virtual memory.

1. Introduction

Although parallel computers were originally designed with processing speed in mind, they have proven equally valuable for their ability to solve problems with very large data requirements. Indeed, parallel computers have opened up a new range of possibilities for scientific computing.

As the capacity of parallel computers has increased, however, so have the appetites of users. Throughout the history of electronic computing, no matter how big and fast the top machines have been, there have always been applications that needed them to be bigger and faster, and it remains true today.

Over thirty years ago, computer architects devised virtual memory to solve this problem for sequential machines [16]. Today's parallel machines typically run traditional sequential virtual memory on the individual nodes. This approach frees the programmer from coding explicit I/O calls, but because it fails to take advantage of aggregate data-parallel operations, it also yields suboptimal I/O performance in *out-of-core* problems, i.e., those whose data re-quirements exceed the size of main memory.

There are multiple paths to reducing I/O times in out-of-core computations. One way is to make each disk access faster; this approach is beyond the scope of the this paper and the ViC* project. Another way is to reduce the number of disk accesses. ViC* is based on this approach.

We know of two ways to reduce the number of disk accesses in an out-of-core data-parallel computation, and the ViC* project uses them both. One is to have the compiler transform the program into one that eliminates many of the disk accesses. The resulting program is essentially the same as the original, but improved. This approach is the focus of this paper.

The other way to reduce the number of disk accesses is to design algorithms that explicitly work with out-of-core data on parallel disks. Since the introduction of the Parallel Disk Model (PDM) by Vitter and Shriver in 1990 [23], there have been significant technical advances on how to carefully plan parallel disk accesses for common data-parallel operations and algorithms [1, 2, 3, 4, 6, 8, 9, 12, 14, 13, 17, 20, 21, 24, 25, 23]. The performance improvements gained by using these methods can be tremendous, and their impacts increase with the problem size. They require a degree of coordination among the processors and disks that unrelated virtual-memory systems on separate nodes cannot provide.

The ViC* approach of built-in virtual-memory support for data-parallel programming allows the memory requirements of application programs to exceed the available memory size without increasing software development time or software complexity. The ViC* compiler transforms the source program to remove many of the parallel disk accesses, and the ViC* runtime system invokes efficient PDM algorithms to perform specific tasks. Programmers do not need specialized knowledge of PDM algorithms in order to avoid huge performance penalties.

To be more specific, the ViC* system is based on using a data-parallel language, in particular C* [22]. The ViC* (<u>Vi</u>rtual-memory <u>C*</u>) compiler transforms a C* program with parallel variables so large that they must reside on disk into a C program with I/O and library calls to access out-

of-core data on a parallel disk system. A ViC* source program does not declare individual variables as disk-resident, or out-of-core; instead, any C* shape may be declared to be `outofcore`, which means that all parallel variables of this shape are out-of-core. The I/O calls added by ViC* read and write sections of out-of-core parallel variables. Before emitting the C code, however, the compiler transforms the program to eliminate many of the I/O calls. The library calls added by ViC* are typically for operations requiring communication in out-of-core parallel variables, e.g., reductions, gets, and sends.

This paper focuses on the ViC* compiler, rather than on the library calls. The library calls are where the efficient PDM algorithms are invoked, and so the library is an important part of the full ViC* system. The compiler, therefore, yields two benefits. First, the transformations it applies directly reduce the number of I/O calls. Second, it makes calls, and enables the programmer to make calls, to the ViC* library, which further reduces the I/O costs.

One principle of this project is to exploit existing languages and software as much as possible. Rather than design a new language, ViC* implements an existing language, C*, with minor extensions. It produces C code, which is processed by host machine compilers.

Why choose C* as a base language? We want an established data-parallel language that is not High-Performance Fortran (HPF). We are interested in data-parallelism because it has proven to be a valuable parallel-programming paradigm and because recent I/O-optimal algorithms fit nicely into it. C* presents different implementation challenges from HPF. In particular, HPF uses arrays, an existing language feature, for parallelism. On the other hand, C* uses shapes, a separate feature not found in sequential C. HPF specifies data distribution at compile time, whereas C* (and ViC*) evaluate shapes at runtime. C* also faces issues of pointer aliasing not permitted in HPF. Many of the issues raised in ViC* implementation are not particular to C*, but are common to compiled data-parallel languages. For example, the language F^{--} [15] bears many similarities to C*. Finally, although many people think of C* solely as a bygone product of Thinking Machines Corporation, there is an active project under the direction of Phil Hatcher at the University of New Hampshire that has produced a C* compiler and runtime system for a distributed-memory model (see [19] and http://www.cs.unh.edu/pjh/cstar/cstar.html).

The remainder of this paper is organized as follows. Section 2 describes virtual memory and its implementation on parallel disk systems. Section 3 presents a brief overview of the C* language and the ViC* extensions. Section 4 discusses program transformations to improve access to out-of-core data. Section 5 describes the runtime interface used to access out-of-core data. Section 6 presents performance measurements. We conclude in Section 7. The full paper [7] describes C* in greater detail and provides details that are omitted here due to space.

2. Virtual Memory

As described by Denning [16] in 1970, *virtual memory* presents the programmer "the *illusion* that he has a very large main memory at his disposal, even though the computer actually has a relatively small main memory." Demand paging is a common implementation of virtual memory, but, as we are about to see, for large data sets there are more efficient alternatives.

Demand paging

Demand paging is a runtime-only mechanism, implemented in the operating system with architectural support and requiring no language or compiler support. Pages are loaded into main memory on demand, i.e., when they are accessed. The program cannot proceed until the data become available. Demand paging services page faults one at a time, based on accesses in a sequential program. It is well suited to a multiprogramming environment that emphasizes throughput rather than latency, since when one process blocks while waiting for page-fault service, another process can run.

Traditional demand paging has relatively poor performance when several passes are made over the same out-of-core data. A typical demand pager replaces the least recently used page (or at least a not very recently used page) with new pages. By the time the last page of out-of-core data is loaded, the first page has been replaced. The result is that each page must be reloaded for each pass through the data.

Just as some optimizing compilers do for in-core data that does not fit in cache, an out-of-core computation can be restructured to combine multiple passes. Even with restructuring to combine multiple passes, with demand paging the speed of the remaining passes is limited by access time to the swap area.

The swap area is typically a partition of a single disk. The access time can be reduced by using a parallel disk system to increase the data transfer rate. Even with the highly unusual configuration of a parallel disk system for swap space, only the transfer rate improves. I/O latency for demand paging does not improve.

I/O latency can be hidden in many out-of-core computations by prefetching and post-writing (see [11] for an example). How good would a demand paging system with a restructuring compiler, parallel disk system, and prefetching/post-writing be? If all computations made sequential passes over the data, it would be quite good.

However, some asymptotically optimal out-of-core algorithms for the Parallel Disk Model (e.g., those for sorting [4, 20, 21, 23], structured permutations [13, 24], and FFTs

[9, 12, 14]) do not access out-of-core data in a simple, sequential fashion. They read and write whole disk blocks, but the blocks may be scattered throughout the parallel disk system. These algorithms require the ability to independently access individual disks. Without an explicit I/O interface, they cannot take full advantage of a parallel disk system. To our knowledge, no demand paging system provides such control.

The principal advantage of demand paging is transparency, not performance. The program takes no part in the managing the virtual memory; indeed it is typically unable to determine that there is virtual memory. For small working sets, demand paging delivers adequate performance.

Virtual memory with ViC*

To support large out-of-core working sets, ViC* modifies the program to explicitly manage its own virtual memory. This is accomplished with minimal changes in the program source, however, requiring only the addition of `outofcore` specifications.

ViC* implements virtual memory for parallel data only; traditional virtual memory mechanisms are adequate to handle instructions and scalar data. Our approach is based on a combination of language features, compiler, and runtime support. We do not require, but can take advantage of, operating system and architectural support for parallel disk systems. Our assumption is that parallel data sets are large enough to warrant special treatment in software. Each access deals with a large amount of data. In fact, the transfer size of a ViC* access is typically much larger than in a traditional demand paging system. Each ViC* access gets at least as much data per disk as in demand paging (if not more), and each ViC* access may be to multiple disks. Consequently, the access cost in ViC* is spread over many more elements than in traditional demand paging.

In this paper we concentrate on data-parallel operations, where the access pattern is the same for all elements. For such code the compiler is able to exploit its knowledge of program structure to reorder accesses, reducing the number of page transfers to and from main memory.

As mentioned previously, ViC* also includes a library of optimal out-of-core algorithms for permutations and other data movement. These algorithms take advantage of an independent I/O interface to a parallel disk system. Such algorithms offer large speedups over conventional in-core algorithms under demand paging. In one case, an explicit out-of-core FFT algorithm was over 144 times faster than a demand-paged version of the traditional in-core Cooley-Tukey method [12].

3. Background concepts and overview of ViC*

This section introduces the parallel programming model and the language features of C* and ViC* that implement it. More information about the C* language appears in [22].

C*, and hence ViC*, supports *data-parallel programming*, in which a sequential program operates on parallel data distributed among a set of *positions*. A *virtual processor* operates on parallel data at each position. The underlying computer multiplexes a set of physical processors among the virtual processors. Scalar data remains global to all virtual processors. This model of programming is also known as SPMD, for Single Program, Multiple Data, a more loosely synchronized software implementation of the SIMD (Single Instruction, Multiple Data) model.

Each parallel variable in C* has a rectangular *shape*, which describes the logical structure of positions. At any point in the program, a *current shape* is in force. *Elemental* parallel operations operate elementwise on data of the current shape. A `with` statement selects the current shape, which is denoted by the reserved word `current`.

All C* operations are controlled by a *context*, which describes the *active positions* in parallel variables of the current shape—those whose virtual processors execute parallel operations. A `where` statement narrows the context, like a parallel `if` statement, by selecting as active a subset of the active positions within the shape. An `everywhere` statement makes all positions active. Exiting a `where` or `everywhere` statement restores the context in force before the statement. Functions inherit the current context and the current shape from their caller.

Parallel *communication* transfers parallel data among the virtual processors. There are several forms of parallel communication. The only one used in this paper is *reduction*, which combines elements of a parallel variable into a scalar result, such as summing the elements.

C* is based on a distributed-memory model of parallel data, with data spread across separate address spaces. In contrast to C arrays, the address of a position in a C* parallel variable is not denotable, and hence individual positions of parallel variables cannot be addressed with pointers. All communication among positions takes place through the explicit communication operations. ViC* implements virtual memory for parallel data by exploiting the distributed-memory model to place out-of-core data on disks and load it into memory as needed.

An Example

The sample ViC* program in Figure 1 illustrates the use of `outofcore` data. The example computes a partial harmonic series for the first $N - 1$ terms, or $(1/k)$ for $k = 1, 2, \ldots, N - 1$, passes its address to an external function, and sums the result. Although the example does not nec-

essarily demonstrate the most efficient means of computing this series, it illustrates a number of issues and optimizations for processing out-of-core data.

The function `harmonize()` declares an out-of-core shape `series` with 2^{30} positions stored on disk. The `with` statement establishes `series` as the current shape. Since `series` is a new shape, the associated context is `everywhere`.

An index variable, `k`, is initialized with the `pcoord()` parallel intrinsic function. The call to `pcoord()` returns the index set along a dimension (in this case, 0) of the current shape. Here it returns `0` in position 0, `1` in position 1, through $2^{30} - 1$ in the last position. The `where` statement narrows the context to positions 1 through $2^{30} - 1$. The parallel variable `harm` is assigned the reciprocals of `k`, i.e., $1/i$ in each position i. Having narrowed the context avoids a division by zero in this expression.

The example calls an external function, `filter()`, with a pointer to `harm`. Function `filter()` also inherits the current shape and context. Such an external call limits the scope of optimizations as described below. On return, a sum-reduction, denoted by the overloaded `+=` operator in C*, returns the sum of the resulting elements of `harm`.

4. Loop Transformations

In this section we discuss transformations of parallel loops—loop fusion, rematerialization and dead store elimination, and scalarization—that improve their performance on out-of-core data. We illustrate their effect on the example from Figure 1 and compare the page I/O counts for parallel data.

In Section 3 we presented an example in the C* language. In this section we describe program transformations in the compiler's intermediate language, using its constructs for parallel data access. In a Section 5 we will discuss the compiler's C language output.

The compiler's `PASS` construct describes a parallel loop over each position, much like Fortran's `FOREACH` loop. Within a `PASS` loop, the `ELT` construct selects the element at the current position. The `PCOORD()` construct evaluates the `pcoord` instrinsic at current position. The global pointer `CONTEXT` points to the current context; a null pointer indicates an `everywhere` context.

Figure 2 shows a C code schema that implements the statements inside the `with` statement of function `harmonize()`. The `PASS` constructs loop for each position. Inside some `PASS` loops, `if` statement restricts execution to the active positions.

The second loop computes the `where` context into an auxiliary parallel variable, `where_1`, which is used in subsequent loops. The global `CONTEXT` is stacked in `context_0` and popped at the end of the `where`

loop	k	where_1	harm	total
1	1024K			1024K
2	1024K	128K		1152K
3	1024K	128K	1024K	2176K
4		128K	1024K	1152K
program	3072K	384K	2048K	5504K

Table 1. Page I/O counts for the program schema in Figure 2 with 8-KB pages.

statement; since functions inherit context when called, `filter()` will also reference `CONTEXT`.

In the fourth loop each processor accumulates a sum in the variable `red_2`. The `sum_reduce()` construct combines the partial sums from each processor and distributes the result to all processors.

Consider the behavior of the schema in Figure 2 with large data sets in a traditional sequential demand-paged environment. Each `PASS` loop will typically page in all of its parallel data operands and write back all parallel results. Consider a single-processor system, such as a DEC Alpha running OSF/1, where an 8-KB page holds 1K `doubles` or `longs`, or 8K booleans (implemented as bytes). We consider only paging due to parallel data, and we assume that the current shape is large enough that no pages are still in memory by the time they are referenced in a subsequent pass. Table 1 shows the expected number of page transfers for parallel variables in each each loop. The total page traffic for this version of the program is 5504K pages. Page prefetching can be used to reduce the latency but does not affect this total I/O count.

Loop Fusion

Loop fusion is a transformation which combines adjacent loops with similar bounds where data dependencies in the loop bodies permit. Loop fusion reduces loop overhead and, more significantly for out-of-core data, improves data locality when the loops access the same data.

A sophisticated C* compiler [19] would fuse these loops, as illustrated in Figure 3. The first three loops of Figure 2 have been fused into a single `PASS`. The remaining loop cannot be fused because of data dependencies: the sum depends on any modifications to elements of `harm` in the call to `filter()`.

In Figure 3, parallel data is reused within the fused loop. All references to `k` are in the first loop, and so this data is traversed exactly once. Variable `where_1` is traversed twice. Table 2 shows the resulting page I/O counts.

```
void  filter(double:current *envelope); /* external computation */

void harmonize() {
  const N = 1 << 30;                    /* N == 2 ** 30 */
  outofcore shape [N]series;            /* series of terms */
  with (series) {                       /* set current shape */
    long int k:current = pcoord(0);     /* index array */
    where (k > 0) {                     /* avoid computing 1/0 */
      double harm:current = 1.0 / k;    /* 1/k (0 < k < N) */
      filter(&harm);                    /* process */
      return += harm;                   /* sum result */
    }
  }
}
```

Figure 1. A sample ViC* program to compute in parallel the terms $1/k$ **for** $0 < k < 2^{30}$ **out of core, call an external function, and return their sum.**

```
bool (*context_0): current = CONTEXT;       /* current context */
bool where_1: current;                      /* context value */
double red_2 = 0;                           /* sum accumulator */

PASS {                                      /* LOOP 1 */
  if (!context_0 || ELT(*context_0))        /* inherited context */
    ELT(k) = PCOORD(0);                     /* set k */
}

PASS {                                      /* LOOP 2 */
  ELT(where_1) = (!context_0 || ELT(*context_0)) && (ELT(k) > 0);
}

CONTEXT = &where_1;                         /* push context */

PASS {                                      /* LOOP 3 */
  if (ELT(where_1))                         /* in new context */
    ELT(harm) = 1.0 / ELT(k);               /* set envelope */
}

filter(&harm);                              /* call */

PASS {                                      /* LOOP 4 */
  if (ELT(where_1))                         /* in new context */
    red_2 += ELT(harm);                     /* sum envelope */
}

red_2 = sum_reduce(red_2);                  /* combine sum */

CONTEXT = context_0;                        /* pop context */

return red_2;
```

Figure 2. C schema for loops in `harmonize()` **in Figure 1. The** `PASS` **construct iterates over all positions, and** `ELT` **evaluates the current position. The variable** `context_0` **stacks the initial context. Temporary variables** `where_1` **and** `red_2` **store the context and summation, respectively.**

Rematerialization

Rematerialization [5] is a transformation that recomputes a variable instead of using its stored value—the inverse of common subexpression elimination. In other compilers, rematerialization typically reduces the number of memory accesses. ViC* uses rematerialization to reduce the number of out-of-core reads at the cost of additional computation. In most current architectures, the cost of disk I/O greatly outweighs the cost of computation. Figure 4 shows the effect of rematerialization on the example.

Rematerialization can be applied when a variable's value is computable from in-core data. Parallel values computed

```
bool (*context_0): current = CONTEXT;        /* inherited context */
bool where_1: current;                       /* context value */
double red_2 = 0;                            /* sum accumulator */

PASS {                                       /* LOOP 1' */
  if (!context_0 || ELT(*context_0))         /* inherited context */
    ELT(k) = PCOORD(0);                      /* set k */
  ELT(where_1) = (!context_0 || ELT(*context_0)) && (ELT(k) > 0);
  if (ELT(where_1)) {                        /* in new context */
    ELT(harm) = 1.0 / ELT(k);                /* set envelope */
  }
}

CONTEXT = &where_1;                          /* push context */

filter(&harm);                               /* call */

PASS {                                       /* LOOP 2' */
  if (ELT(where_1))                          /* in new context */
    red_2 += ELT(harm);                      /* sum envelope */
}

red_2 = sum_reduce(red_2);                   /* combine sum */

CONTEXT = context_0;                         /* pop context */

return red_2;
```

Figure 3. C schema for `harmonize()` **in Figure 1 with the first four loops of Figure 2 fused into a single loop.**

```
bool (*context_0): current = CONTEXT;        /* inherited context */
bool where_1: current;                       /* context value */
double red_2 = 0;                            /* sum accumulator */

PASS {                                       /* LOOP 1'' */
  if (!context_0 || ELT(*context_0))         /* inherited context */
    ELT(k) = PCOORD(0);                      /* set k */
  ELT(where_1) = (!context_0 || ELT(*context_0)) && (ELT(k) > 0);
  if (ELT(where_1)) {                        /* in new context */
    ELT(harm) = 1.0 / ELT(k);                /* set envelope */
  }
}

CONTEXT = &where_1;                          /* push context */

filter(&harm);                               /* call */

PASS {                                       /* LOOP 2'' */
  if (!context_0 || ELT(*context_0))         /* inherited context */
    ELT(k) = PCOORD(0);                      /* rematerialize new context */
  ELT(where_1) = (!context_0 || ELT(*context_0)) && (ELT(k) > 0);
  if (ELT(where_1)) {
    red_2 += ELT(harm);                      /* sum envelope */
  }
}

red_2 = sum_reduce(red_2);                   /* combine sum */

CONTEXT = context_0;                         /* pop context */

return red_2;
```

Figure 4. Rematerialization. The values of `k` **and** `where_1` **are recomputed rather than being read from storage.**

loop	k	where_1	harm	total
$1'$	1024K	128K	1024K	2176K
$2'$		128K	1024K	1152K
program	1024K	256K	2048K	3328K

Table 2. Page I/O counts for the program schema in Figure 3 after loop fusion.

loop	k	where_1	harm	total
$1''$	1024K	128K	1024K	2176K
$2''$			1024K	1024K
program	1024K	128K	2048K	3200K

Table 3. Page I/O counts for the program schema in Figure 4 after rematerialization.

loop	k	where_1	harm	total
$1'''$		128K	1024K	1152K
$2'''$			1024K	1024K
program		128K	1024K	2176K

Table 4. Page I/O counts for the program schema in Figure 4 after dead store elimination.

from scalars and the pcoord() function are suitable for rematerialization, as are variables computed from other values already available in a loop. A context which may be everywhere, such as the one established by the with statement, is also a good candidate for rematerialization, since it may require no data. In the second loop of Figure 4 the context value, where_1 is rematerialized from the initial context and the expressions pcoord(0) and k > 0. The resulting page I/O counts are shown in Table 3. Total page traffic is reduced to 3200K pages.

Dead Store Elimination

Dead store elimination is applied after rematerialization has eliminated references to stored parallel variables. Because the local parallel variable k is rematerialized in the second loop, its value is never read from disk. Consequently, the assignment to k is *dead*, and so its computed vales need not be written to disk. On the other hand, the current context in where_1 is implicitly passed to filter(), so it must be written.

Determining that a parallel assignment is dead requires dataflow analysis of variables shared between loops. Dead store elimination does not alter the loop structure, but it eliminates the need to write back the out-of-core data. The resulting page I/O counts are shown in Table 4. The total

page I/O count has been reduced to less than half the original count.

Scalarization

Scalarization replaces parallel variables with scalars. It applies to intermediate parallel variables that are used only within a loop and neither read nor written to storage. Scalarization reduces the overhead of selecting a position as well as the memory requirements in a loop. Scalarization is applied after rematerialization has eliminated reads and dead store elimination has eliminated writes.

In Figure 5, this transformation is applied to the local variable k. In the second PASS, the current context is also replaced with a scalar. Only harm and where_1 in the first loop and harm in the second loop remain as out-of-core variables. These are exactly the variables visible to the function filter().

Scalarization requires dataflow analysis of variables shared within a loop. Although scalarization does not affect the page I/O requirements, it improves in-core access within the PASS loops, an important consideration for in-core shapes.

Summary

For this example, the loop transformations described above reduce page I/O counts by 60%. Experienced programmers often apply such optimizations, particularly when I/O is explicit. As can be seen by comparing Figures 1 and 5, the transformations tend to obscure the program structure. Adding I/O calls results in a program more like Figure 6. This sequence of transformations is the sort of programming process that is better left to a compiler such as ViC*.

5. Runtime Support

In this section, we describe the ViC* runtime I/O interface for parallel data loops. After the loops have been restructured, the ViC* compiler inserts I/O calls and expands parallel variable references.

Figure 6 shows the final C code of harmonize() after applying all the transformations in Section 4. Each PASS loop is expanded into an outer *sectioning* loop and an inner in-core vector loop. References to parallel variables are replaced by references to in-core strips. Other calls to runtime functions set up I/O and manage loop iteration.

Each loop begins with a call to ACCESS_DATA() for each parallel variable used in the loop. This function opens file descriptors and allocates in-core buffers for out-of-core data. Access is given as read-only ('R'), write-only ('W'), modify (read and write, 'M'), or read-context ('C'). The ACCESS_DATA() function detects aliasing among parallel

```
bool (*context_0): current = CONTEXT;        /* inherited context */
bool where_1: current;                        /* context value */
double red_2 = 0;                             /* sum accumulator */

PASS {                                        /* LOOP 1''' */
  long k_3;                                   /* scalar ELT(k) */
  if (!context_0 || ELT(*context_0))          /* inherited context */
    k_3 = PCOORD(0);                          /* set k */
  ELT(where_1) = (!context_0 || ELT(*context_0)) && (k_3 > 0);
  if (ELT(where_1)) {                         /* in new context */
    ELT(harm) = 1.0 / k_3;                    /* set envelope */
  }
}

CONTEXT = &where_1;                           /* push context */

filter(&harm);                                /* call */

PASS {                                        /* LOOP 2''' */
  long k_4;                                   /* scalar ELT(k) */
  bool where_5;                               /* scalar ELT(where_1) */
  if (!context_0 || ELT(*context_0))          /* inherited context */
    k_4 = PCOORD(0);                          /* rematerialize new context */
  where_5 = (!context_0 || ELT(*context_0)) && (k_4 > 0);
  if (where_5) {
    red_2 += ELT(harm);                       /* sum envelope */
  }
}

red_2 = sum_reduce(red_2);                    /* combine sum */

CONTEXT = context_0;                          /* pop context */

return red_2;
```

Figure 5. Scalarization. Parallel variable k and, in the second loop, the current context are replaced by scalar temporaries.

data references and combines the references. It also verifies that the parallel data has the current shape.

At first glance, it might appear that reading a context is no different than reading any other parallel bool. It turns out that there are two reasons to treat the context specially. First, it provides a simple programming check. Recall that the pointer to a context may be null, which would indicate an everywhere context. In such a case, no read would occur. For all other out-of-core parallel data, the pointer must be non-null; an access type other than read-context with a null data pointer triggers a runtime error. (The origin of such an error would be the ViC* compiler rather than the programmer.) Second, treating the context specially enables a runtime optimization. Again supposing that the context pointer is null, consider what happens when we perform write-only access on some other out-of-core parallel variable a. Because the context is everywhere, there is no need to first read a; every position will be written. On the other hand, when the context is not known to be everywhere, we must first read a in order to maintain values in the inactive positions. When ACCESS_DATA() finds that the context is everywhere, it can automatically optimize access to write-only data within the loop.

Within the loop, all data movement is managed by ITERATE_STRIPS(), which returns a true value as long as there are additional parallel data to process. For in-core shapes no I/O is required, and ITERATE_STRIPS() returns true only once. For out-of-core shapes, ITERATE_STRIPS() also manages I/O, prefetching the next strip and writing back previous strips. I/O is striped across disks in a parallel disk system. Asynchronous I/O is overlapped with in-core computation.

Once data is in-core, INCORE_STRIP() locates the current in-core strip of each parallel variable. All these strips have the same number of in-core positions, as returned by COUNT_STRIP(), regardless of the element size. Thus, for example, there are as many in-core elements of harm as of context_0 in the second loop of Figure 6, although elements of harm occupy eight bytes and elements of context_0 occupy one byte.

The inner loop processes in-core data, iterating a virtual-processor index through the in-core positions. Within the inner loop, all parallel data references are replaced by strip array references, indexed by the virtual processor number.

```
bool (*context_0): current = CONTEXT;               /* inherited context */
bool where_1: current;                               /* context value */
double red_2 = 0;                                    /* sum accumulator */

ACCESS_DATA(context_0,'C')                           /* loop over context */
ACCESS_DATA(&where_1,'W')                            /* loop over where_1 */
ACCESS_DATA(&harm,'W')                               /* loop over harm */
while (ITERATE_STRIPS()) {                            /* LOOP 1'''' */
  int vp_10;                                          /* VP index */
  bool *strip_11 = (bool*)INCORE_STRIP(context_0);   /* in-core data */
  bool *strip_12 = (bool*)INCORE_STRIP(&where_1);
  double *strip_13 = (double*)INCORE_STRIP(&harm);
  for (vp_10 = COUNT_STRIP(); 0<=--vp_10; ITERATE_PCOORD()) {
    long k_3;                                         /* scalar ELT(k) */
    if (!context_0 || strip_11[vp_10])               /* inherited context */
      k_3 = PCOORD(0);                                /* set k */
    strip_12[vp_10] = (!context_0 || strip_11[vp_10]) && (k_3 > 0);
    if (strip_12[vp_10]) {                            /* in new context */
      strip_13[vp_10] = 1.0 / k_3;                    /* set envelope */
    }
  }
}

CONTEXT = &where_1;                                  /* push context */

filter(&harm);                                       /* call */

ACCESS_DATA(context_0,'C')                           /* loop over context */
ACCESS_DATA(&harm,'R')                               /* loop over harm */
while (ITERATE_STRIPS()) {                            /* LOOP 2'''' */
  int vp_13;                                          /* VP index */
  bool *strip_14 = (bool*)INCORE_STRIP(context_0);   /* in-core data */
  double *strip_15 = (double*)INCORE_STRIP(&harm);
  for (vp_13 = COUNT_STRIP(); 0<=--vp_13; ITERATE_PCOORD()) {
    long k_4;                                         /* scalar ELT(k) */
    bool where_5;                                     /* scalar ELT(where_1) */
    if (!context_0 || strip_14[vp_13])
      k_4 = PCOORD(0);
    where_5 = (!context_0 || strip_14[vp_13]) && (k_4 > 0);
    if (where_5) {
      red_2 += strip_15[vp_13];                       /* sum envelope */
    }
  }
}

red_2 = sum_reduce(red_2);                           /* combine sum */

CONTEXT = context_0;                                 /* pop context */

return red_2;
```

Figure 6. Expanded loop code. `PASS` **constructs have been expanded into inner in-core and outer sectioning loops. Parallel data is prefetched and accessed through in-core strip pointers.**

If necessary, each iteration calls `ITERATE_PCOORD()` to set the next next value of `PCOORD(0)`. The resulting code structure encourages optimization by the final C compiler.

6. Performance

We measured performance of the loops in Figure 6, varying the number of positions in the shape `series` and characteristics of parallel data access. The test system is a DEC 2100 server with two 175-MHz Alpha processors, eight disks, and 320 MB of main memory. The main computation uses a single thread. An I/O thread manages I/O for each of the disks. Each processor can run any ready thread.

After the transformations described above, the first loop writes 9 bytes of data and the second loop reads 8 bytes for each position in the shape `series`. For this example, the main memory capacity is a little over 35 million positions. Out-of-core data is accessed through the file system by way of a 64 MB in-core buffer pool.

Figure 7 shows the average time per position for problem sizes ranging from 1 million to 100 million positions. An in-core shape relies on demand paging to access the data. Out-of-core shapes manage I/O through the file system. The standard ViC* implementation uses asynchronous I/O, but a synchronous implementation reduces the memory buffer requirements. The number of parallel disks varied among 1, 2, 4, and 8.
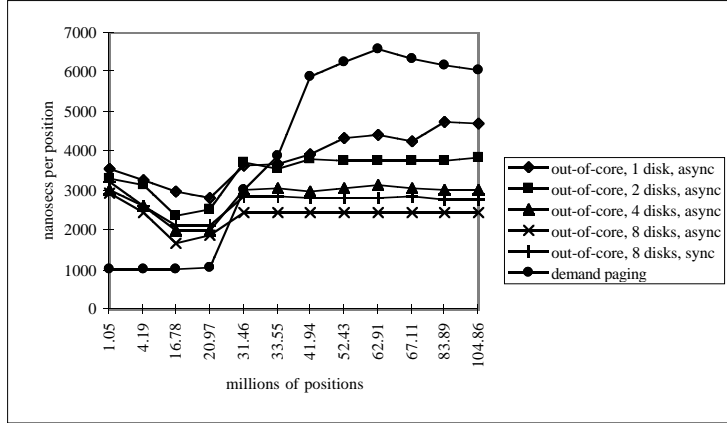
Figure 7. Performance of the loops in Figure 6 with in-core and out-of-core data implementations. (Note that the horizontal axis is not linear.)

In-core data access with demand paging outperforms out-of-core data access at small problem sizes, but it degrades rapidly at large problem sizes. Below 30 million positions, there is little paging, and the time to process each position averages around one microsecond. As the problem size approaches the main memory size, demand paging degrades performance sharply. For problem sizes above 50 million positions, demand paging adds about 5 microseconds to the processing time for each position. An address space of 1 GB limits in-core problem sizes to 100 million positions.

Although out-of-core data access shows higher overhead costs at small problem sizes, it maintains performance at large problem sizes. At small problem sizes, fixed-cost loop setup operations dominate. The system buffer cache makes a copy of out-of-core data as it is read or written, thus adding a constant to the time for each position. The buffer cache is effective in reducing I/O below 30 million positions, the region where demand paging is also effective, but at a higher cost. As the problem size becomes greater than the size of main memory, out-of-core data-access times remain stable. At 100 million positions, out-of-core data access across eight disks takes between 2 and 3 microseconds per position.

At small problem sizes an in-core `shape` yields better performance than an `outofcore shape`. The programmer should choose between them based on available memory size and expected data accesses, although the the runtime system could be extended to make such decisions automatically using an adaptive algorithm.

Even with a single disk, the ViC* out-of-core code outperforms demand paging at large problem sizes. Increasing the number of parallel disks further reduces the access time. Synchronous I/O across eight disks reduces performance somewhat by removing overlap between computation and

I/O, but the ViC* code continues to access disks in parallel. The buffer pool is large enough that the increased buffer sizes available for synchronous I/O provide little benefit.

A programmer can achieve similar performance manually by using a similar strategy. The ViC* library uses standard I/O calls, with the addition of a small overhead. The advantage of the ViC* I/O library lies in its use of high-performance interfaces. Such I/O interfaces calls complicate the program structure, as seen in Figure 6. Further, some I/O optimizations rely on data-flow assumptions that must be re-validated when the program is modified. The ViC* compiler provides a way to do this, and it allows the programmer to easily convert between in-core and out-of-core programming. Of course, a skilled and patient programmer may achieve better performance by global program restructuring.

7. Conclusion

We have described an implementation of virtual memory for out-of-core data-parallel programming with a parallel disk system. ViC* divides responsibility for memory management between the programmer, who declares `outofcore` shapes, the compiler, which restructures parallel operations, and the runtime system, which manages I/O data transfers and buffering. This approach differs from conventional demand-paged virtual memory, which operates at the level of instructions and memory references. The ViC* runtime requires minimal operating system support, exploiting parallel and asynchronous I/O where that is available.

By using explicit runtime I/O interfaces, the ViC* compiler is able to manage parallel virtual memory at a higher level, coordinating the out-of-core data transfers across an entire loop. The compiler also emits calls to a runtime library that invokes efficient parallel disk algorithms for out-

of-core communication functions. In both ways, the ViC* system significantly reduces disk-access costs for out-of-core data.

Acknowledgments

References

[1] A. Aggarwal and C. G. Plaxton. Optimal parallel sorting in multi-level storage. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 659–668, Jan. 1994.

[2] L. Arge. The buffer tree: A new technique for optimal I/O-algorithms. In *4th International Workshop on Algorithms and Data Structures (WADS)*, pages 334–345, Aug. 1995.

[3] L. Arge, D. E. Vengroff, and J. S. Vitter. External-memory algorithms for processing line segments in geographic information systems. In P. Spirakis, editor, *Proceedings of the Third Annual European Symposium on Algorithms (ESA '95)*, volume 979 of *Lecture Notes in Computer Science*, pages 295–310. Springer-Verlag, Sept. 1995.

[4] R. D. Barve, E. F. Grove, and J. S. Vitter. Simple randomized mergesort on parallel disks. *Parallel Computing*, 23(4–5):601–631, June 1997.

[5] P. Briggs, K. D. Cooper, and L. Torczon. Rematerialization. In *Proceedings of the ACM SIGPLAN '92 Conference on Programming Language Design and Implementation*, pages 311–321, June 1992.

[6] Y.-J. Chiang, M. T. Goodrich, E. F. Grove, R. Tamassia, D. E. Vengroff, and J. S. Vitter. External-memory graph algorithms. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 139–149, Jan. 1995.

[7] A. Colvin and T. H. Cormen. ViC*: A compiler for virtual-memory C*. Technical Report PCS-TR97-323, Dartmouth College Department of Computer Science, Nov. 1997.

[8] T. H. Cormen. Fast permuting in disk arrays. *Journal of Parallel and Distributed Computing*, 17(1–2):41–57, Jan. and Feb. 1993.

[9] T. H. Cormen. Determining an out-of-core FFT decomposition strategy for parallel disks by dynamic programming. Technical Report PCS-TR97-322, Dartmouth College Department of Computer Science, July 1997. To appear in [18].

[10] T. H. Cormen and A. Colvin. ViC*: A preprocessor for virtual-memory C*. Technical Report PCS-TR94-243, Dartmouth College Department of Computer Science, Nov. 1994.

[11] T. H. Cormen and M. Hirschl. Early experiences in evaluating the Parallel Disk Model with the ViC* implementation. *Parallel Computing*, 23(4–5):571–600, June 1997.

[12] T. H. Cormen and D. M. Nicol. Performing out-of-core FFTs on parallel disk systems. Technical Report PCS-TR96-294, Dartmouth College Department of Computer Science, Aug. 1996. To appear in *Parallel Computing*.

[13] T. H. Cormen, T. Sundquist, and L. F. Wisniewski. Asymptotically tight bounds for performing BMMC permutations on parallel disk systems. Technical Report PCS-TR94-223, Dartmouth College Department of Computer Science, July 1994. Preliminary version appeared in *Proceedings of the 5th Annual ACM Symposium on Parallel Algorithms and Architectures*. Revised version to appear in *SIAM Journal on Computing*.

[14] T. H. Cormen, J. Wegmann, and D. M. Nicol. Multiprocessor out-of-core FFTs with distributed memory and parallel disks. In *Proceedings of the Fifth Workshop on I/O in Parallel and Distributed Systems (IOPADS '97)*, pages 68–78, Nov. 1997. Also Dartmouth College Computer Science Technical Report PCS-TR97-303.

[15] B. D. de Dinechin, G. Elsesser, G. Fischer, B. Johnson, T. MacDonald, R. W. Numrich, and J. L. Steidel. Definition of the F^{--} extension to Fortran 90. In *Proceedings of the 10th International Workshop on Languages and Compilers for Parallel Computing*, Aug. 1997. To appear.

[16] P. J. Denning. Virtual memory. *ACM Comput. Surv.*, 2(3):153–189, Sept. 1970.

[17] M. T. Goodrich, J.-J. Tsay, D. E. Vengroff, and J. S. Vitter. External-memory computational geometry. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 714–723, Nov. 1993.

[18] *Proceedings of the Workshop on Algorithms for Parallel Machines*, 1996–97 Special Year on Mathematics of High Performance Computing, Institute for Mathematics and Its Applications, University of Minnesota, Minneapolis, Sept. 1996.

[19] A. J. Lapadula and K. P. Herold. A retargetable C* compiler and run-time library for mesh-connected MIMD multicomputers. Technical Report TR 92-15, University of New Hampshire, 1992. Revised by Phil Hatcher, October 1993.

[20] M. H. Nodine and J. S. Vitter. Deterministic distribution sort in shared and distributed memory multiprocessors. In *Proceedings of the 5th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 120–129, June 1993.

[21] M. H. Nodine and J. S. Vitter. Greed sort: Optimal deterministic sorting on parallel disks. *J. ACM*, 42(4):919–933, July 1995.

[22] Thinking Machines Corporation. *C* Programming Guide*, May 1993.

[23] J. S. Vitter and E. A. M. Shriver. Algorithms for parallel memory I: Two-level memories. *Algorithmica*, 12(2/3):110–147, Aug. and Sept. 1994.

[24] L. F. Wisniewski. *Efficient Design and Implementation of Permutation Algorithms on the Memory Hierarchy*. PhD thesis, Dartmouth College Department of Computer Science, Mar. 1996.

[25] D. Womble, D. Greenberg, S. Wheat, and R. Riesen. Beyond core: Making parallel computer I/O practical. In *DAGS '93*, June 1993.