

Programming many-core architectures - a case study: dense matrix computations on the Intel SCC processor

Bryan Marker¹, Ernie Chan¹, Jack Poulson², Robert van de Geijn¹, Rob F. Van der Wijngaart³, Timothy G. Mattson⁴, Theodore E. Kubaska⁵

¹ Dept. of Computer Science, The Univ. of Texas at Austin, Austin, Texas 78712. ² Institute for Computational Engineering and Sciences, The Univ. of Texas at Austin, Austin, Texas 78712. ³ Intel Corporation, Santa Clara, California 95054. ⁴ Intel Corporation, DuPont, Washington 98327. ⁵ Intel Corporation, Hillsboro, Oregon 97124.

SUMMARY

A message passing, distributed-memory parallel computer on a chip is one possible design for future, many-core architectures. We discuss initial experiences with the Intel Single-chip Cloud Computer research processor, which is a prototype architecture that incorporates 48 cores on a single die that can communicate via a small, shared, on-die buffer. The experiment is to port a state-of-the-art, distributed-memory, dense matrix library, Elemental, to this architecture and gain insight from the experience. We show that programmability addressed by this library, especially the proper abstraction for collective communication, greatly aids the porting effort. This enables us to support a wide range of functionality with limited changes to the library code. Copyright © 2011 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: collective communication, dense linear algebra library, many-core architecture

1. INTRODUCTION

The computer industry is at a crossroads. The number of transistors on a chip continues to climb with successive generations of process technology (Moore's law) while the power available to a socket is decreasing. This has led to a "power wall" and has shifted the focus of computer architecture from raw performance to performance per watt.

A well-known response to the power wall problem is to replace complex cores running at high frequencies with multiple simple but low power cores within a chip [6]. The major microprocessor vendors currently offer CPUs with modest numbers of cores (two to eight) organized around a cache-coherent shared address space. These multicore processors in many ways appear to the programmer as a familiar multiprocessor, multi-socket system integrated onto a single chip. Cache-coherent shared memory is convenient for the programmer since the hardware creates the illusion of a single, coherent address space that spans multiple cores and maintains consistency on behalf of the programmer. But this abstraction adds overhead that grows with the number of cores and hence may not be scalable to support large numbers of cores. An alternative approach is to model chips with multiple cores after clusters, which are parallel architectures with scalable disjoint memories that lack cache coherence. An instance of this approach is the Intel Single-chip Cloud Computer (SCC).

*Correspondence to: Dept. of Computer Science, The Univ. of Texas at Austin, Austin, Texas 78712.
bamarker@cs.utexas.edu

The SCC processor [17, 20] is a 48-core “concept vehicle” created by Intel Labs as a platform for many-core software research. The chip presents to the programmer a collection of cores with private memories, connected by an on-die network. These can be programmed as a “cluster on a chip” with messages moving around the network to coordinate execution of processes running on the cores and communicate data between those processes. In addition to this logically distributed memory, the SCC processor has two shared address spaces: one on-die and one off-chip. Neither of these address spaces provides any level of cache coherence between cores, which makes the chip highly scalable but leaves the burden of maintaining a consistent view of these address spaces to the programmer.

In this paper, we describe the results of an effort to port a major software library, the Elemental library [21] for dense matrix computations on distributed-memory computer architectures, to this platform. To do so, we start with a minimal programming environment, RCCE [20, 28], that consists of synchronous point-to-point communication primitives. This communication layer allows all issues related to coherency to be hidden in the passing of messages, at the expense of placing the entire burden of parallelization on the library programmer. We show that by adding a few commonly used and one novel collective communication to this layer, the entire Elemental library, which has functionality similar to ScaLAPACK [7] and PLAPACK [26], is successfully ported with relatively little effort. The conclusions we draw from this experience are:

- Message passing can be an effective way to avoid having to provide cache coherence in many-core architectures.
- Software that can be cast in terms of interleaved stages of computation and structured communication, namely collective communication, can be supported by distributed-memory, many-core architectures such as SCC. One collective communication not commonly contained in other message passing libraries, which we call *permutation* (see Section 4.3), was discovered in the process and added to our set of supported collectives. Its utility extends well beyond Elemental, enabling many advanced parallelization strategies.
- Focus on programmability in Elemental greatly aided our porting efforts. Here, programmability means layering code that re-uses abstractions for communication, computation, etc. throughout the library. Collective communication functionality is contained within a layer on which higher-level code and algorithms are built using a well-defined API. This made porting easier because the collective communication layer could be modified independently of all dependent code as long as the expected functionality was maintained.

The rest of the paper is organized as follows. We provide a brief overview of SCC in Section 2. Section 3 discusses Elemental along with our efforts to port it to SCC. The port to SCC is discussed in Section 4. Performance results are provided in Section 5. An explanation of why this approach was undertaken instead of alternatives is given in Section 6. We discuss the lessons learned from our efforts and plans for future work in Section 7.

2. THE SCC PROCESSOR

The Single-chip Cloud Computer (SCC) processor is an experimental processor [17] from Intel Labs. It uses an architecture that can scale to many hundreds of cores as the density of transistors that can be placed within a single chip continues to increase. The SCC project explores hardware questions such as low-power routers, explicit power management, and scalable network-on-a-chip architectures. Its most important role, however, is as a platform for many-core software research. Here, we explore the programmability and scalability features of such a chip.

The SCC processor was created through a software/hardware co-design process. As the processor was designed, the RCCE native message passing environment was developed for the chip [20]. By using a functional emulator, we were able to develop applications and propose changes to the processor architecture as it was being developed. In this section, we briefly review the architecture of the SCC processor and RCCE.

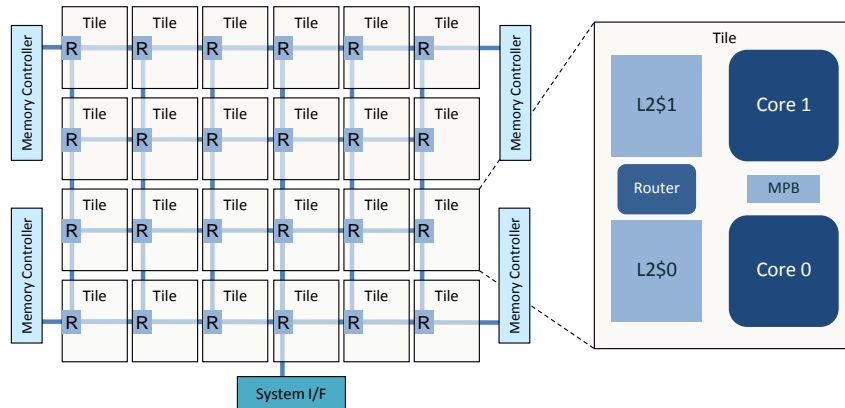


Figure 1. SCC architecture is comprised of a 4×6 grid of tiles where each tile contains a pair of cores (with L1 and L2 caches), a router, and 16 KB of shared SRAM to serve as a message passing buffer.

2.1. The SCC Architecture

The SCC processor architecture is shown in Figure 1. The processor consists of 24 tiles organized into a 4×6 grid. The routers (denoted by R in the figure) implement fixed X-Y routing. The chip has 16 to 64 GB of off-die DRAM memory.

A tile contains a pair of minimally modified P54C processor cores [1], each with an independent L1 (16 KB data and 16 KB instruction) and L2 (256 KB unified instruction/data) cache. The cores are second generation Pentium® processors selected because of their low-power, in-order architecture. The choice of this core seriously limits the raw performance of the chip but does not impede progress on the key research vectors for the project, e.g., programmability, scalability, power management.

Each tile also includes a router and a 16 KB block of SRAM. These memory blocks are organized into a shared address space visible to all cores on the chip. This memory was added to support the movement of L1 cache lines between cores and hence is called the “message passing buffer” (MPB). It is important to appreciate that the processor does not maintain cache coherency between cores for any memory region on the chip. All issues of coherency or consistency are managed explicitly by the programmer. When working with the SCC processors, programmers are exposed to three distinct address spaces:

- A private memory in off-chip DRAM for each core. This memory is coherent with an individual cores’s L1 and L2 caches.
- The MPB that has 24×16 KB of shared memory in SRAM.
- A shared-memory off-chip address space in DRAM.

The MPB is an important feature of the SCC processor. Since the private memory associated with each core is a distinct address space, cores cannot exchange information by passing pointers. The MPB lets cores exchange information in the form of messages. Because it is on-die, the MPB provides a low-overhead mechanism to move blocks of data from one core’s L1 to another’s L1 and ultimately between private memories. It would be possible to exchange data through the shared-memory in the off-chip DRAM, but this would suffer from higher latency and lower bandwidth.

In this paper, we view the SCC processor as a collection of cores with local memories, communicating through a message passing library described below, and we test its programmability as an integrated cluster. Future work by our group and others in the SCC research community will explore other programming models, in particular models that make direct use of the shared, off-chip memory available on the SCC processor.

2.2. RCCE Communication Library

RCCE (pronounced “rocky”) [20, 28] is a light-weight communication library developed by Intel for the SCC processor. It defines low-latency mechanisms to move data stored in the private memory of one core to the private memory of another core. The most common usage model for RCCE assumes synchronous, two-sided communication. Cores that need to exchange information wait for all participating cores to reach corresponding points in their execution. Then they cooperatively exchange data as needed. This approach is common with Message Passing Interface (MPI) [14] applications targeting cluster computers.

At the lowest level, RCCE provides a one-sided communication layer. The basic RCCE design treats the MPB as a set of 8 KB buffers, each designated to one core. To move a cache line from one core to another, the sending core “puts” (copies) a cache line into its own buffer from which the receiving core “gets” (copies) the cache line, thereby moving it into its own L1 cache. Programmers need to coordinate movement of cache lines into and out of the MPB. This is done with “flags”, i.e., synchronization variables within RCCE.

The basic one-sided communication API within RCCE is flexible and can handle a wide range of communication patterns, but it can be a complicated approach. We quickly recognized that higher level, two-sided communication primitives, much like the `send` and `receive` functions that MPI provides, were needed. Several more simplified, MPI-like functions were added on an as-needed basis, but not asynchronous communication.

The exclusion of asynchronous communication in RCCE deserves an explanation. The SCC processor typically operates with a Linux kernel running on each core. Given Linux, we can execute with multiple threads on each core, thereby supporting asynchronous communication conveniently. An alternative mode of using SCC, however, uses a low level operating system-less mode, which we call *baremetal mode*. We designed RCCE so programs can be built and executed in Linux and baremetal mode without a change in source code. The cost is that programmers must convert asynchronous algorithms to ones that use synchronous communication, and that the potential for overlap of computation and communication is removed. However, the SCC architecture itself, lacking out of order execution and a message co-processor, does not allow such overlap anyway. Moreover, there are costs associated with supporting asynchronous communications, specifically with respect to message buffering, traversing and maintaining message queues, and allocating and maintaining synchronization flags. We also note that the SCC does not experience the high inter-process latencies typical of clusters, reducing the potential savings of overlap. Hence, the greatest cost to the programmer is the recoding required to avoid deadlock, which can often arise with strictly synchronous messaging. An important observation of this paper is that for the dense linear algebra functions we have explored, the restriction of synchronous communication is not a problem because of the way the ported library is programmed.

3. ELEMENTAL

In this section, we give a brief overview of the Elemental library. Cholesky factorization is used as a representative operation to illustrate some of the programming issues and how Elemental addresses them.

3.1. Why Elemental?

LINPACK [8] can be considered the first numerical package that tried to address programmability in addition to functionality and numerical stability. Its developers adopted the Basic Linear Algebra Subprograms (BLAS) [19] interface for portable performance. Subsequently, the Linear Algebra PACKage (LAPACK) [2] was developed to provide higher performance on cache-based architectures by adopting new layers of the BLAS [9, 10] as well as added functionality and stability. As distributed-memory architectures became increasingly common, a level of abstraction was needed to support dense linear algebra on these systems. This led to ScaLAPACK [7], which extended LAPACK functionality to distributed-memory computer architectures.

Applications		
Elemental Solvers		
Elemental BLAS/Decomposition/Reduction/...		
Elemental Local Operations	Elemental Redistribution Operations	
Local Compute Kernels (BLAS/LAPACK)	Packing Routines	Collective Communication (MPI or RCCE)

Figure 2. Layering of the Elemental library.

The goal of ScaLAPACK was performance and functionality and on those fronts it has been a very successful package. However, the project did not place emphasis on programmability. PLAPACK [26], a dense linear algebra package for distributed-memory machines similar in functionality to ScaLAPACK, added programmability as a key focus to its design. The central idea is that the algorithms used for dense matrix computations should be apparent in the source code. The group behind PLAPACK has continued this line of research, exploring systematic derivation of linear algebra algorithms supported by clear abstractions to support their expression in source code. This led to a new, sequential dense linear algebra library, `libflame` [29], and more recently a new dense linear algebra library for distributed-memory architectures, Elemental [21]. It is this library that is at the heart of our experiment since it was designed for conventional clusters but appeared suitable to be ported to architectures like SCC.

Elemental solves scalability problems encountered in PLAPACK and programmability problems encountered in ScaLAPACK, which is explained in Section 6.2. For SCC Elemental has a few obvious advantages: 1) it uses a simple data distribution; 2) it is carefully layered, as shown in Figure 2; 3) it uses abstractions that allow the programmer to code at the level at which one reasons about the algorithm; and, importantly, 4) all communication is cast in term of collective communication. These features improve programmability of the library and greatly eased the porting effort.

3.2. A Motivating Example: Cholesky Factorization

If $A \in \mathbb{R}^{n \times n}$ is a symmetric, positive-definite matrix, then Cholesky factorization computes $A \rightarrow U^T U$ where U is an upper triangular matrix. One algorithm, often called the right-looking variant, is presented using FLAME notation [4, 15] in Figure 3. This figure presents a blocked algorithm that casts most computation in terms of matrix-matrix computations (level-3 BLAS), allowing it to attain high performance on cache-based architectures.

The FLAME notation helps solve the programmability problem by providing a methodology that constructively derives algorithms to be correct [3], meaning that as functionality is added to libraries like Elemental and `libflame`, a high level of confidence can be placed in the correctness of the resulting algorithms.

3.3. Representing Algorithms in Code

The first feature of Elemental that aids programmability is that it is coded in C++ using modern, object-oriented coding practices, a deviation from the implementation of the alternative packages ScaLAPACK and PLAPACK, which were respectively coded in Fortran77 and C.

The Elemental implementation of the Cholesky algorithm is given in Figure 4. This code is representative of the layer that is labeled “BLAS/Decomposition/Reduction/...” in Figure 2. Those familiar with the FLAME project will recognize that, like other FLAME related APIs, the code resembles the algorithm in Figure 3, hiding indexing details. In the case of Elemental, the code also encapsulates details about how the matrix is distributed among cores (or, in the case of MPI,

<p>Algorithm: $A := \text{CHOL_BLK}(A)$</p> <p>Partition $A \rightarrow \left(\begin{array}{c c} A_{TL} & A_{TR} \\ \hline * & A_{BR} \end{array} \right)$</p> <p style="padding-left: 20px;">where A_{TL} is 0×0</p> <p>while $m(A_{TL}) < m(A)$ do</p> <p style="padding-left: 20px;">Determine block size b</p> <p style="padding-left: 20px;">Repartition</p> <p style="padding-left: 40px;">$\left(\begin{array}{c c} A_{TL} & A_{TR} \\ \hline * & A_{BR} \end{array} \right) \rightarrow \left(\begin{array}{c c c} A_{00} & A_{01} & A_{02} \\ \hline * & A_{11} & A_{12} \\ \hline * & * & A_{22} \end{array} \right)$</p> <p style="padding-left: 40px;">where A_{11} is $b \times b$</p> <hr style="width: 50%; margin-left: 40px;"/> <p style="padding-left: 40px;">$A_{11} := \text{CHOL}(A_{11})$</p> <p style="padding-left: 40px;">$A_{12} := A_{11}^{-H} A_{12}$ (TRSM)</p> <p style="padding-left: 40px;">$A_{22} := A_{22} - A_{12}^H A_{12}$ (TRIANGULAR RANK-K)</p> <hr style="width: 50%; margin-left: 40px;"/> <p style="padding-left: 20px;">Continue with</p> <p style="padding-left: 40px;">$\left(\begin{array}{c c} A_{TL} & A_{TR} \\ \hline * & A_{BR} \end{array} \right) \leftarrow \left(\begin{array}{c c c} A_{00} & A_{01} & A_{02} \\ \hline * & A_{11} & A_{12} \\ \hline * & * & A_{22} \end{array} \right)$</p> <p>endwhile</p>
--

Figure 3. Blocked, right-looking algorithm for computing the Cholesky factorization. Note that the algorithm is for both a real and complex valued matrix A where A^H denotes conjugate transposition.

processes). Commonly used code patterns and functionality are abstracted (e.g. in functions), and higher-level code is layered on top. By re-using this code, Elemental partially solves the programmability problem by using an API that reduces the opportunity for introducing “bugs” when the algorithm in Figure 3 is translated into the code in Figure 4.

3.4. Distribution and Redistribution

To understand how parallelism is expressed in the Elemental code, one must first understand a little about the data distributions used by Elemental and how redistribution is expressed.

When parallelizing a sub-operation on distributed-memory architectures, data is initially distributed in some specified fashion among the processes. Ideally, each process has all the data it needs, so all can simultaneously perform independent local operations. In practice this typically requires data to be duplicated or redistributed among the processes before local computation commences. Often, the local computation is a contribution to a global result, and data must be reduced (e.g., summed) leaving it in some prescribed distribution. Redistribution and/or reduction requires communication.

To support scalability, dense linear algebra libraries distribute matrices by viewing the processes as a logical two-dimensional mesh. These libraries attain load balance as the computation proceeds by wrapping matrices cyclically around the process mesh [16, 24, 25]. Elemental partially solves the programmability problem by choosing the simplest such distribution: the p processes are viewed as forming an $r \times c$ logical mesh, and the elements of a given matrix A are wrapped using an *elemental* 2D cyclic distribution, which means that element (i, j) is assigned to process $(i\%r, j\%c)$. This is in contrast to PLAPACK and ScaLAPACK, which use a more complex block cyclic distribution. In those packages, the blocksize b_{distr} is used. Blocks of size b_{distr} by $r \times b_{\text{distr}}$ in PLAPACK and of size b_{distr} by b_{distr} in ScaLAPACK are wrapped around the processes grid in a 2D cyclic fashion. As a result, indexing and redistribution are more complicated because the “owning” process for element (i, j) is not simply $(i\%r, j\%c)$ as it is in Elemental. Thus, the code within Elemental related to the distribution and redistribution of data is much simpler than in the other packages. Furthermore, PLAPACK’s distribution, tied to the number of rows in the process grid, makes code mildly non-scalable when the number of processors becomes large enough and the matrix fills all available memory. In [21] it is shown that Elemental’s simplification does *not* adversely affect performance on traditional clusters. The benefit of this simplification is that it makes the communication (redistribution) layer easier to adapt to new platforms.

```

1  template<typename T> void
2  elemental::lapack::internal::Chol_blk
3  ( DistMatrix<T,MC,MR>& A )
4  {
5      const Grid& g = A.GetGrid();
6
7      DistMatrix<T,MC,MR>
8          ATL(g), ATR(g), A00(g), A01(g), A02(g),
9          ABL(g), ABR(g), A10(g), A11(g), A12(g),
10         A20(g), A21(g), A22(g);
11
12     DistMatrix<T,Star,Star> A11_Star_Star(g);
13     DistMatrix<T,Star,VR > A12_Star_VR(g);
14     DistMatrix<T,Star,MC > A12_Star_MC(g);
15     DistMatrix<T,Star,MR > A12_Star_MR(g);
16
17     PartitionDownDiagonal
18     ( A, ATL, ATR,
19       ABL, ABR, 0 );
20     while( ABR.Height() > 0 )
21     {
22         RepartitionDownDiagonal
23         ( ATL, /**/ ATR,  A00, /**/ A01, A02,
24           /*****/ /*****/
25             /**/ A10, /**/ A11, A12,
26           ABL, /**/ ABR,  A20, /**/ A21, A22 );
27
28         A12_Star_MC.AlignWith( A22 );
29         A12_Star_MR.AlignWith( A22 );
30         A12_Star_VR.AlignWith( A22 );
31         //-----//
32         A11_Star_Star = A11;
33         lapack::internal::LocalChol( Upper,
34                                     A11_Star_Star );
35         A11 = A11_Star_Star;
36
37         A12_Star_VR = A12;
38         blas::internal::LocalTrsm
39         ( Left, Upper, ConjugateTranspose, NonUnit,
40           (T)1, A11_Star_Star, A12_Star_VR );
41
42         A12_Star_MC = A12_Star_VR;
43         A12_Star_MR = A12_Star_VR;
44         blas::internal::LocalTriangularRankK
45         ( Upper, ConjugateTranspose,
46           (T)-1, A12_Star_MC, A12_Star_MR,
47           (T)1, A22 );
48         A12 = A12_Star_MR;
49         //-----//
50         A12_Star_MC.FreeAlignments();
51         A12_Star_MR.FreeAlignments();
52         A12_Star_VR.FreeAlignments();
53
54         SlidePartitionDownDiagonal
55         ( ATL, /**/ ATR,  A00, A01, /**/ A02,
56           /**/ A10, A11, /**/ A12,
57           /*****/ /*****/
58           ABL, /**/ ABR,  A20, A21, /**/ A22 );
59     }
60     return;
61 }

```

Figure 4. Elemental implementation of the blocked, right-looking algorithm for the Cholesky factorization.

ScaLAPACK and PLAPACK cast most communication in terms of collective communications within rows and/or columns of processes. Elemental takes this one step further and casts *all* communication in terms of collective communication. This is important, since tracking down all cases of point-to-point communication to ensure that they can be implemented in terms of the synchronous communications supported by RCCE is laborious and error-prone. The individual collective communications used by Elemental are isolated in a specific layer of the library and are themselves easy to cast in terms of synchronous point-to-point communications, as we discuss next.

4. RETARGETING TO SCC

At the onset of this research, there were multiple reasons why retargeting Elemental to SCC appeared to be a natural fit. Since the primary purpose of SCC is to investigate programmability issues related to many-core architectures, it made sense to tap into the focus of the FLAME project on programmability. Elemental, which builds on the FLAME project, was developed for conventional clusters and one of the memory models of SCC allows us to view it as a distributed-memory system (cluster) on a single chip. Elemental uses collective communication, which maps well to the synchronous point-to-point communication supported by RCCE. Message passing, used via collective communication in the case of Elemental, is a model that avoids having to explicitly manage coherency between cores since this is handled within the message passing primitives. Finally, the FLAME project's emphasis on program correctness and the abstractions developed for Elemental gave a high degree of confidence in that code base, meaning that as the port proceeded there was never a question of whether there was a latent logic error in Elemental. If something did not work, the problem was always with the small number of routines that were tailored to SCC and RCCE.

4.1. Changes

To understand where changes had to be made requires an explanation of what happens when a redistribution is triggered by a command like the one in line 32 of Figure 4. The assignment is overloaded to redistribute data [21]. In this case, the processes recognize the “before” and “after” distributions and determine that data from all processes must be communicated to all processes by an allgather. Before a routine that implements allgather can be called, though, the local data must be rearranged (packed) into a convenient format. After completion of the collective communication, it must again be locally rearranged (unpacked) by each process. In between is a call to an MPI collective communication on a conventional cluster. This call needed only to be replaced by a call to an equivalent RCCE collective communication. Thus, it is only in the box labeled “Collective Communication” in Figure 2 that changes were made. Said another way, the focus on programmability in designing and layering Elemental allowed us to retarget the library to a new architecture with minimal changes outside that layer in the code because mostly functions in that layer were adapted and all other code called those functions.

4.2. Preliminary work

When this research commenced, no SCC processor was available for the port, so it was performed with the aid of the previously mentioned RCCE emulator. The major challenge was that only some of the collective communications used by Elemental were part of RCCE's collective communication library. We discuss how we dealt with that challenge in the next section. Conveniently, at the heart of the FLAME project is a methodology for systematically deriving different algorithmic variants for a given linear algebra operations [27]. Different variants often require different redistributions and hence even when initially only a few collective communications were available in RCCE, broad functionality of Elemental came on-line. The benefit of eventually having all collective communication used in Elemental available was that all variants for all operations supported by

Before			After		
Allgather					
p_0	p_1	p_2	p_0	p_1	p_2
x_0			x_0	x_0	x_0
	x_1		x_1	x_1	x_1
		x_2	x_2	x_2	x_2
Alltoall					
p_0	p_1	p_2	p_0	p_1	p_2
$x_0^{(0)}$	$x_0^{(1)}$	$x_0^{(2)}$	$x_0^{(0)}$	$x_1^{(0)}$	$x_2^{(0)}$
$x_1^{(0)}$	$x_1^{(1)}$	$x_1^{(2)}$	$x_0^{(1)}$	$x_1^{(1)}$	$x_2^{(1)}$
$x_2^{(0)}$	$x_2^{(1)}$	$x_2^{(2)}$	$x_0^{(2)}$	$x_1^{(2)}$	$x_2^{(2)}$

Figure 5. Illustrations of allgather and alltoall with three processes.

allgather	alltoall
$l = (me - 1 + p) \% p$ $r = (me + 1) \% p$ $i = me$ for $j = 1, \dots, p - 1$ $k = (i + 1) \% p$ Send x_i to p_l Receive x_k from p_r $i = k$ endfor	for $j = 0, \dots, p - 1$ $i = (j - me + p) \% p$ if $i \neq me$ Send x_i to p_i Receive x_i from p_i endif endfor

Figure 6. Cyclic algorithm for allgather and pairwise exchange algorithm for alltoall.

Elemental were also available, meaning that the best-performing algorithmic variant for distributed-memory computing could be employed. For example, there are three commonly used variants for Cholesky factorization. The initial port of Elemental to SCC only supported one of those variants, the left-looking variant, because RCCE does not implement all of the communication patterns in MPI or even all of those used by Elemental. Eventually, the right-looking variant in Figures 3 and 4, which parallelizes more naturally, was supported as we introduced more collective communication operations.

4.3. Collective communication

All of the communication between processes in Elemental is cast in terms of collective communication. When this research commenced, RCCE only provides a small set of unoptimized collective communication routines (collectives) such as broadcast, where a vector of data on one process is sent to all other processes. Elemental uses several standard collective communication routines provided by MPI that are not supported by RCCE, so we implemented these operations using only the simple, synchronous point-to-point communication routines (`send` and `receive`). The interface to the collective communication routines purposely mimics those of corresponding MPI routines.

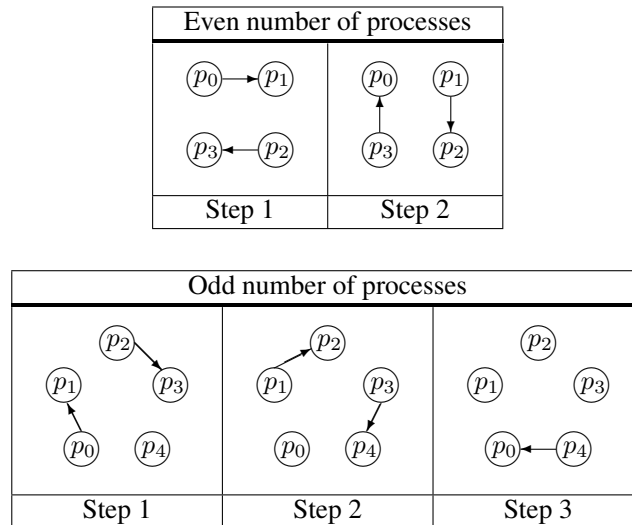


Figure 7. Deadlock prevention within a ring communication pattern when given even and odd number of processes.

Allgather For the `allgather` operation, assume a vector of data, x , is partitioned into subvectors x_i , and initially each process p_i owns only x_i . Upon completion each of the p processes owns the entire vector x .

A relatively simple algorithm for `allgather` is the cyclic, or bucket, algorithm [5] is given in Figure 6 (left), where me denotes the rank of the calling process. Within an iteration of this algorithm, each process sends its local contribution to its neighboring process on the (logical) left and receives from the (logical) right. This algorithm inherently sends data in a circular communication pattern. Since both `send` and `receive` in RCCE are blocking, deadlock occurs with this cyclic algorithm if implemented in a single communication step. All processes would first call the `send` routine and will block until the corresponding `receive` is posted, but no process will do so since all are blocked on the `send`.

This deadlock can be easily resolved using two steps for an even number of processes. In the first step, all the even numbered processes send while all odd numbered processes receive. In the second step, evens receive and odds send. A problem arises when the number of processes is odd, but we can avoid deadlock by introducing a third step. In the first step, all evens send to odds, excluding the wrap-around where process p_{p-1} sends to p_0 . In the second step, all odds send to evens. Finally, the wrap-around occurs where the first and last ranked processes communicate with each other. As a result, deadlock can be avoided by detecting a ring communication pattern and performing this odd and even decomposition of the sends and receives. We illustrate deadlock prevention of these two cases in Figure 7 using four and five processes as examples.

Alltoall The `alltoall` operation performs a permutation of the vector x on each process. Initially, each process p_i owns a vector $x^{(i)}$ that is partitioned where $x_j^{(i)}$, $j = 0, \dots, p-1$. Upon completion the process p_j owns the permuted vector $x_j^{(i)}$, $i = 0, \dots, p-1$.

A relatively simple algorithm for `alltoall` is a staged pairwise exchange algorithm [23] is given in Figure 6 (right). This algorithm is deadlock free because only distinct pairs of cores communicate with each other during each stage. We can play the trick of having the lower ranked process send first while the higher ranked process receives, and then reverse roles. The algorithm can be improved slightly for even numbers of processes [28], yielding a provably optimal schedule.

Send-receive (permutation) On the surface, one important redistribution in Elemental is not a collective communication. Instead, it requires point-to-point communication similar to MPI's

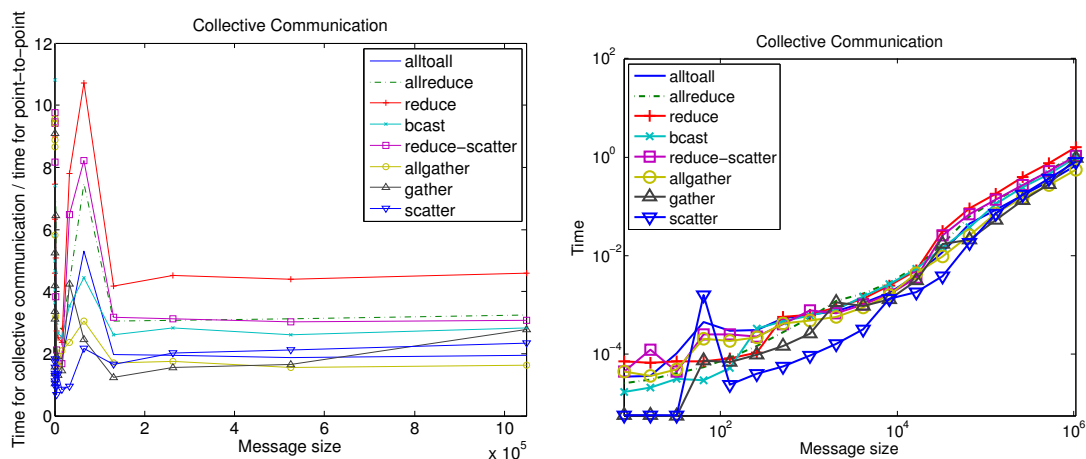


Figure 8. Collective communication runtimes on 48 cores of SCC. On the left is performance relative to point-to-point communication and on the right is runtime.

`sendreceive`. The MPI routine can be thought of as combining `send` and `receive` (possibly from a different process) into a single routine. In general this would be difficult to implement in terms of synchronous point-to-point communications. Fortunately, the communication that requires this in Elemental, and many other codes, can be viewed as a collective communication that implements a permutation where every process sends data to one other process and also receives data from one other process. It is this permutation functionality that we support.

We implement this routine by first distributing all the sending and receiving ranks for each process via an `alltoall` operation, so all processes know how all other processes will communicate. The resulting communication graph consists of a collection of linear (open) chains and/or cycles, each of which can always be implemented in a maximum of three communication stages (two in the case of a linear chain or cycle with an even number of nodes). We can “cache” this communication pattern so that the communication graph construction is only performed once and is subsequently reused in Elemental. As a result, we do not need to call `alltoall` each time the permutation is invoked.

5. PERFORMANCE

We provide performance data for the collective communication routines described above, which are used in Elemental on SCC. In Figure 8, we show the performance of the routines run on 48 cores for varying message sizes. On the left the runtime of the collective is scaled by the runtime of point-to-point communication. On the right is the runtime for each operation. Since the cost of point-to-point is effectively a lower bound for the cost of the collective communications, we find that the simple implementations discussed in the last section are quite effective.

We now discuss Elemental performance, not only in terms of efficient utilization and scalability, but also in terms of the effort required when porting the software. We compare the performance on varying numbers of cores for three dense matrix operations: Cholesky factorization; LU factorization with partial pivoting, $PA := LU$ where $A \in \mathbb{R}^{n \times n}$, L is a lower triangular matrix, U is an upper triangular matrix, and P is a permutation matrix; and general matrix-matrix multiplication (GEMM), $C := \alpha AB + \beta C$ where $A \in \mathbb{R}^{m \times k}$, $B \in \mathbb{R}^{k \times n}$, and $C \in \mathbb{R}^{m \times n}$.

We tuned the algorithmic block size and the grid configuration within Elemental. We linked Elemental with the Intel Math Kernel Library (MKL) 8.1.1 for the computational kernels on each core using double precision floating point real arithmetic. We compare the scalability of Elemental versus a sequential MKL implementation: `dpotrf` for CHOL, `dgetrf` for LU, and `dgemm` for GEMM. MKL is used for single-core performance.

Notice that using the default clock frequency of 533 MHz, each core has a theoretical peak performance of 533 MFLOPS giving SCC a total theoretical peak performance of 25.584 GFLOPS. However, this is not a realistic performance target given that MKL's `dgemm` on a single core only achieved around 120 MFLOPS. Hence, this is not an experiment on how to achieve near-peak performance.

We provide performance results in Figures 9, 10, and 11. In the left of those figures, we compare the scalability of Elemental by fixing a few sample problem sizes and showing performance using varying numbers of cores. Linear speedup is shown relative to the sequential MKL implementation run on one core. On the right, we show the breakdown of the different component costs within implementations using all 48 cores and varying the problem sizes. The computation cost only entails the time each process spent executing a sequential kernel when linked to MKL. The communication cost is the time spent in collective communication routines, which includes copying between private memory and MPB. The overhead contains all remaining components of the execution such as packing and unpacking data to and from contiguous application-level buffers.

Consider the performance of matrix multiplication for a single core, found on the left of Figure 11, which used the sequential MKL `dgemm` kernel.[†] Given the cache-friendliness of GEMM, an ambitious speedup on n cores would be n times this single-core performance. As we increased the number of cores available to Elemental, performance increased with no obvious “knees” in the curves indicating diminishing marginal utility. Given the way Elemental distributes data and parallelizes algorithms, we believe it would scale well on processors similar to SCC with even more cores. Such scalability is seen in [21] on cluster computers composed of many cores.

Notice that scalability improved with larger problem sizes as the computational time on p processes is $O(n^3/p)$ while the communication and packing related time is $O(n^2/\sqrt{p})$ (and hence $O(n^3)$ versus $O(n^2)$ when p is fixed). This trend continued with larger problem sizes than those shown. This is typical behavior for these operations on clusters, as costly communication is a larger portion of execution time for smaller problem sizes than larger ones. Notice the decrease of the communication cost portion in the component graphs as the problem sizes increase, and the computational portion simultaneously increases. As communication and overhead costs are relatively smaller portions of overall performance for larger problem sizes, speedup improves as the problem size increases.

In Figure 12 (left), we show the performance of a representative set of operations supported in Elemental that were ported to SCC. These include all level-3 BLAS and several LAPACK-level operations. This graph illustrates the breadth of functionality that was quickly ported due to the focus on programmability. For each of these operations, Elemental contains multiple algorithmic variants thanks to this focus on programmability. Different variants exhibit contrasting performance characteristics and use different communication patterns. Without the communication routines described in Section 4, only two of the operations in this graph worked on SCC. If only a single variant of each operation were available, we could not test correctness of early ports of Elemental as easily because we would have to implement new algorithm variants in addition to porting the library. Instead, we were able to choose variants of those operations that only called collective communication functions available in RCCE. With the additional communication routines of Section 4, all variants of the remaining operations worked immediately. We show the default variant of each operation in this graph.

6. ALTERNATIVES

It is hard to appreciate the success of the described approach without putting it in perspective by discussing alternatives. To this end, we now briefly discuss the level of effort that was required to port the LINPACK benchmark, the effort that we believe would have been required to port ScaLAPACK, and an ongoing effort to port a multicore solution that uses shared memory.

[†]Due to per-core memory limits, only two of the problem sizes fit on a single core.

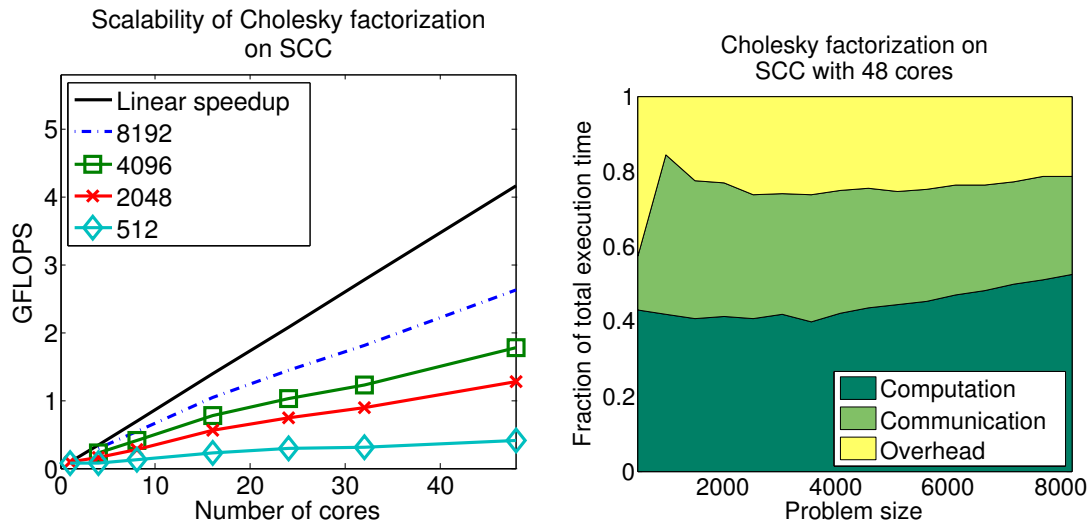


Figure 9. Scalability (left) and cost breakdown for 48 cores (right) of Elemental's implementation of Cholesky factorization.

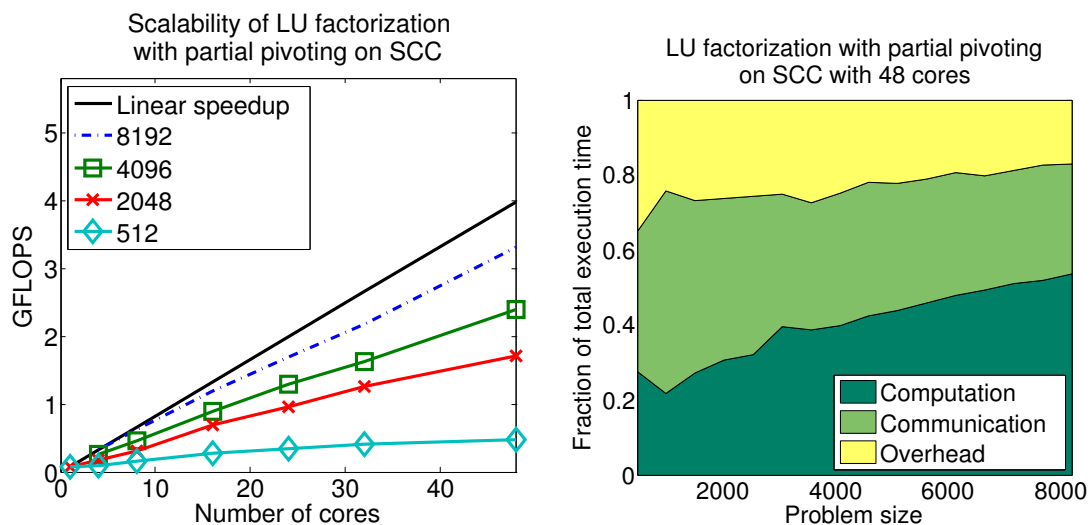


Figure 10. Scalability (left) and cost breakdown for 48 cores (right) of Elemental's implementation of LU factorization with partial pivoting.

6.1. High-Performance LINPACK

The High-Performance LINPACK (HPL) is a highly specialized implementation of the LINPACK benchmark [11] for massively-parallel, distributed-memory systems that was partially ported to SCC by one of the co-authors. HPL performs a large LU factorization with partial pivoting, and much like ScaLAPACK, it uses a block cyclic data distribution and fundamentally does not address programmability. Details such as pipelining where communication and computations are overlapped are exposed directly within the code.

In order to port HPL to SCC, we replaced all asynchronous MPI communication calls with synchronous RCCE routines. Deadlocks had to be detected and explicitly avoided. This required non-trivial analysis of the HPL communication patterns underlying the point-to-point messages, which greatly complicated the port.

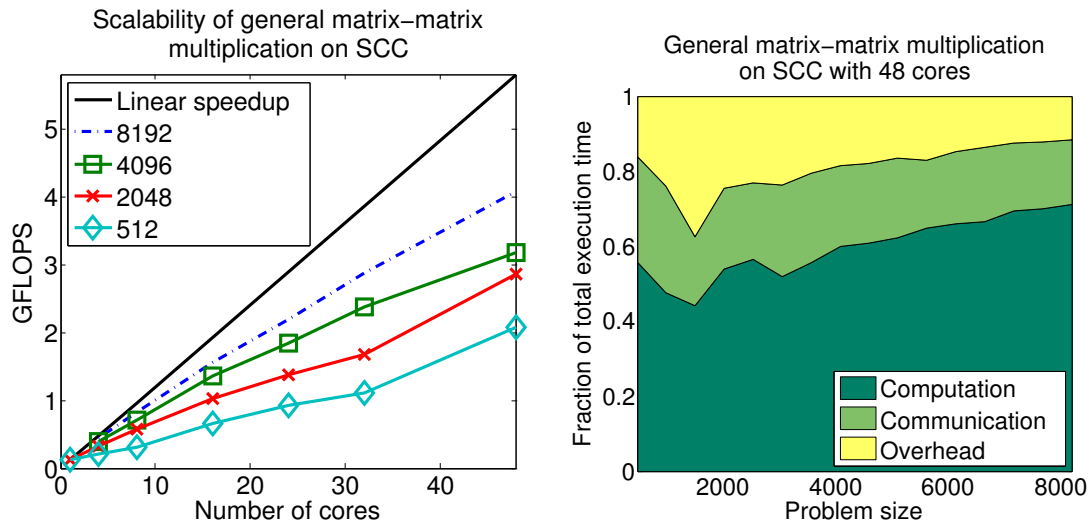


Figure 11. Scalability (left) and cost breakdown for 48 cores (right) of Elemental's implementation of general matrix-matrix multiplication where the matrix dimensions are $m = n$ and $k = 1280$.

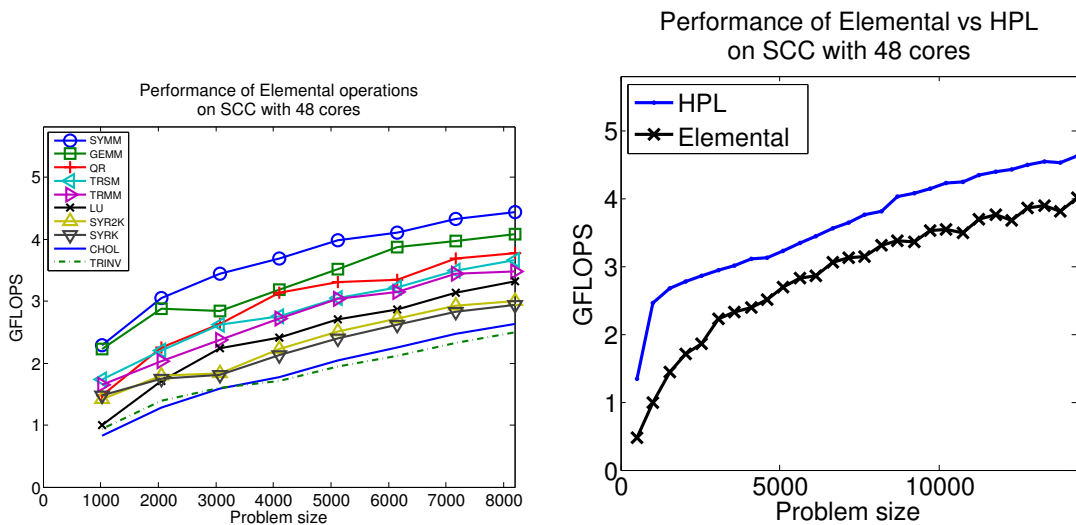


Figure 12. Performance of a cross-section of operations supported by Elemental that have been ported to SCC (left) and LU factorization with partial pivoting for Elemental versus HPL (right) using all 48 cores.

Elemental's implementation of LU factorization with partial pivoting is compared against HPL in Figure 12 (right). HPL required tuning of a number of parameters. We were not able to tune all parameters because HPL was not completely ported to SCC. The tuning parameters, which often involve algorithmic variants, create a large quantity of code to port, and the effort required cannot be justified merely for the purpose of comparison. We tuned as much as possible for a single problem size and used the best choices for all runs with the exception of the block size, for which the optimal setting changed for small and large problems.

Although Elemental's performance is lower than HPL's, the difference is fairly modest, especially for larger problem sizes, and is expected to narrow even further as Elemental is further optimized. We are already investigating improvements to Elemental's LU factorization implementation and describe some promising optimizations to the port in Section 7.2. Moreover, Elemental is a general-purpose library created to enable many algorithms to be developed for distributed-memory

computers whereas HPL is a benchmark meant solely to achieve good performance for this one particular operation. We consider Elemental's relative performance deficit the result of a reasonable compromise between speed and versatility/programmability.

Regarding the latter, we note that even the incomplete port of HPL to RCCE required the assistance of the author of the original code because of the complicating effect of the point-to-point communications. Elemental, in contrast, was much easier to port completely, as it fully isolates the required data transport in a modest collection of generic collective communications [13]. Although the main author of Elemental is a co-author of this paper, the port was accomplished by another co-author of this paper who had little experience with Elemental and distributed-memory computing and received virtually no help from the author of Elemental.

6.2. Porting ScaLAPACK

One could argue that a comparison between Elemental and ScaLAPACK would have been a better experiment. However, ScaLAPACK contains many point-to-point communications and a much larger body of code than HPL, which prevented us from attempting the port.

To quantify this last statement, we point out a few key issues. First, major design decisions regarding ScaLAPACK were made prior to the arrival of MPI. Second, ScaLAPACK, by design, is layered and coded to closely resemble LAPACK. As a result, the library-level code is layered upon the parallel BLAS (PBLAS) layer, which itself is layered upon standard (local) BLAS and the Basic Linear Algebra Communication Subprograms (BLACS), a communication layer that has an interface that resembles the BLAS interface [12]. The BLACS themselves are coded in terms of what, at the time, were a myriad of native communication libraries. The most commonly used implementation is now layered upon MPI. The BLACS include both point-to-point and collective primitives.

In principle the BLACS collective communications should be easy to port. In practice the BLACS implement an array of algorithms for collective communication without relying on the MPI interface. Still, it would be a matter of simplifying this implementation so that they call only the collective communications that were developed as part of our effort. This might possibly go at the expense of performance since ScaLAPACK depends on pipelining between communication and computation in a number of important routines in order to reduce communication overhead.

The more troublesome aspect of a port of ScaLAPACK comes from its use of point-to-point communications. In the PBLAS we found 37 instances of calls to `DGESD2D`, the BLACS `send` primitive for communicating double precision data. At the library level (LAPACK-level functionality), we found 168 such instances. Each of these may need to be examined to determine whether the communication can be performed synchronously and possibly reimplemented so that it can be performed synchronously. Not counted here are a large number of calls in the ScaLAPACK test suite and redistribution routines.

The point is that porting ScaLAPACK is possible but labor intensive. By comparison, the only place where point-to-point communications are called by Elemental is in its communication layer where we automatically avoid deadlock and communication serialization. On conventional architectures, Elemental delivers performance that is competitive with, and often exceeds, that of ScaLAPACK [21].

6.3. Porting a shared-memory solution

Another question is whether the given hardware should be viewed and programmed as a distributed-memory or a shared-memory architecture. When we started our research, we did both. Elemental was ported to examine how easily and effectively the SCC processor could be programmed as a distributed-memory architecture. Our `libflame` library [29] together with the SuperMatrix runtime system [22] was ported to see how easily the system could be programmed as a shared-memory architecture. This approach views matrices as blocks (submatrices) that are units of data and operations with these blocks as units of computation. Execution of a `libflame` routine builds a Directed Acyclic Graph (DAG) of tasks and dependencies between tasks, which is then scheduled at runtime by the SuperMatrix runtime system. This approach has been highly effective on

conventional multicore architectures. Both the shared-memory and distributed-memory approaches were ported easily to the SCC emulator. The Elemental solution was subsequently ported easily to the architecture itself, as described in this paper.

By contrast the SuperMatrix runtime system did not port easily to the SCC architecture. First, handling private versus shared memory required a software layer that turned out to be difficult to develop and debug due to the lack of hardware coherency support. Second, the performance (even after a year of development) continues to lag greatly behind that of the Elemental solution. Finally, calls to the MKL BLAS implementation often crash the architecture, which indicates that shared memory use by these routines is in conflict when running on multiple cores concurrently.

The fact is that by viewing and programming the processor as a distributed-memory architecture and porting a modern distributed-memory dense linear algebra library to it quickly yielded success. We believe this provides evidence that viewing many-core processors as distributed-memory architectures is a plausible solution and makes the case that when experimenting on emerging architectures, embracing software that pays attention to programmability is advantageous.

7. CONCLUSION

In this paper, we have described our experiences related to the porting of a major software library, Elemental, to the SCC research processor. We started with the conjecture that for some problem domains software supported coherency of data on many-core architectures can be achieved by viewing the architecture as a distributed-memory parallel computer architecture and communicating data via message passing constructs. For the domain of dense matrix computations, the results provide early evidence that this is indeed the case when one starts with a library that already targets distributed-memory architectures and is very carefully layered. It is shown that a minimal set of communication primitives is needed to support this, namely collective communication.

7.1. Insights

We targeted a problem domain that is thought to be well-understood but has struggled with the complications of parallel computing for two decades. Fortunately, that struggle allowed insight to be gained from legacy libraries, ScaLAPACK and PLAPACK, yielding a properly layered library, Elemental, that fundamentally addresses the programmability problem for the domain of dense matrix computations. As a result, this library ported easily to SCC processor, building on the RCCE communication library. By building on primitive point-to-point communications, the RCCE library is able to exploit these routines to provide a full set of collective communications.

A question is how representative the domain of dense matrix computations is of other software libraries and “real” applications. A careful look at our results shows that any application that casts its communication in terms of stages of computation interleaved with stages of communication that can be implemented with synchronous communication should port to this kind of platform. One may argue that few applications fall into this category, but notice that one could have come to the same conclusion for the domain of dense matrix computations had one started with ScaLAPACK. Thus, the real story is that by building on prior art like ScaLAPACK and PLAPACK, we managed to effectively layer a new library for the domain of dense matrix computations that had this desired property. Similarly, there are likely other domains that can be recast in such a way. The point is that the arrival of many-core architectures is an opportunity to reexamine and rearchitect existing software.

7.2. Future Directions

Very little effort has been made to optimize the Elemental port to SCC. We would especially like to optimize the bottom layer of Figure 2 to improve performance. Some of the collective communication routines have opportunities for optimization. For example, they view the process grid as a linear array of processes, which it is not. Furthermore, the sequential MKL library is unoptimized for SCC. It should be updated to take advantage of the L1 and L2 caches of the

Pentium processor to improve the base performance for single-core computation. Lastly, the packing operations of Elemental copy data into contiguous memory buffers to call RCCE communication routines, which subsequently copy data from those buffers to the MPB in roughly 8 KB chunks. By breaking this boundary to provide communication routines that skip this intermediate copy, we can substantially reduce the overhead cost seen in Figures 9, 10, 11.

More generally, it is unwise to bet on only one solution, given the uncertainty of future architectures. As part of the FLAME project, a number of solutions have been developed for parallelizing dense linear algebra libraries. First, the sequential `libflame` library can be linked to multithreaded BLAS. Elemental and SuperMatrix, described in Section 6.3, are two other solutions well-suited to be ported to SCC. Another solution views the cores as a distributed-memory architecture and uses message passing to implement the SuperMatrix scheduler and the passing of blocks of data [18]. Interestingly, this approach, which uses message passing and the RCCE communication library, again yielded a relatively simple port. In the future, we intend to compare these alternative approaches to the one presented in this paper.

Additional Information

For additional information on the Formal Linear Algebra Methods Environment (FLAME), visit

<http://www.cs.utexas.edu/users/flame/>.

For further information on Elemental, visit

<http://code.google.com/p/elemental/>.

ACKNOWLEDGEMENTS

We thank the other members of the FLAME team for their support. We also like to thank Jesper Larsson Traff from the University of Vienna for his help with the implementation of `alltoall` and Antoine Petitet from the University of Tennessee for his help with the HPL port.

This work was supported in part by the Intel Labs Academic Research Office, which also donated access to the SCC processor. Bryan Marker is funded by a Sandia National Laboratories fellowship, and Jack Poulson is funded by a Computational Applied Mathematics fellowship from The University of Texas at Austin.

REFERENCES

1. D. Anderson and T. Shanley. *Pentium processor system architecture*. Addison Wesley, 1995.
2. E. Anderson, Z. Bai, J. Demmel, J. E. Dongarra, J. DuCroz, A. Greenbaum, S. Hammarling, A. E. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide*. SIAM, Philadelphia, 1992.
3. P. Bientinesi, J. A. Gunnels, M. E. Myers, E. S. Quintana-Ortí, and R. A. van de Geijn. The science of deriving dense linear algebra algorithms. *ACM Transactions on Mathematical Software*, 31(1):1–26, March 2005.
4. P. Bientinesi, E. S. Quintana-Ortí, and R. A. van de Geijn. Representing linear algebra algorithms in code: The FLAME application programming interfaces. *ACM Trans. Math. Soft.*, 31(1):27–59, March 2005.
5. E. Chan, M. Heimlich, A. Purkayastha, and R. van de Geijn. Collective communication: theory, practice, and experience: Research articles. *Concurr. Comput. : Pract. Exper.*, 19:1749–1783, September 2007.
6. A. P. Chandrakasan, M. Potkonjak, R. Mehra, J. Rabaey, and R. W. Brodersen. Optimizing power using transformations. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 14(1):12–31, January 1995.
7. J. Choi, J. J. Dongarra, R. Pozo, and D. W. Walker. Scalapack: A scalable linear algebra library for distributed memory concurrent computers. In *Proceedings of the Fourth Symposium on the Frontiers of Massively Parallel Computation*, pages 120–127. IEEE Comput. Soc. Press, 1992.
8. J. J. Dongarra, J. R. Bunch, C. B. Moler, and G. W. Stewart. *LINPACK Users' Guide*. SIAM, Philadelphia, 1979.
9. J. J. Dongarra, J. Du Croz, S. Hammarling, and I. Duff. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Soft.*, 16(1):1–17, March 1990.
10. J. J. Dongarra, J. Du Croz, S. Hammarling, and R. J. Hanson. An extended set of FORTRAN basic linear algebra subprograms. *ACM Trans. Math. Soft.*, 14(1):1–17, March 1988.
11. J. J. Dongarra, P. Luszczyk, and A. Petitet. The linpack benchmark: Past, present, and future. concurrency and computation: Practice and experience. *Concurrency and Computation: Practice and Experience*, 15:2003, 2003.
12. J. J. Dongarra, R. A. van de Geijn, and R. C. Whaley. Two dimensional basic linear algebra communication subprograms. In *Proceedings of the Sixth SIAM Conference on Parallel Processing for Scientific Computing*, March 1993.
13. S. Gorbach. Send-recv considered harmful: Myths and realities of message passing. *ACM Transactions on Programming Languages and Systems*, 26, 2004.

14. W. Gropp, E. Lusk, and A. Skjellum. *Using MPI*. The MIT Press, 1994.
15. J. A. Gunnels, F. G. Gustavson, G. M. Henry, and R. A. van de Geijn. Flame: Formal linear algebra methods environment. *ACM Transactions on Mathematical Software*, 27(4):422–455, December 2001.
16. B. A. Hendrickson and D. E. Womble. The torus-wrap mapping for dense matrix calculations on massively parallel computers. *SIAM J. Sci. Stat. Comput.*, 15(5):1201–1226, 1994.
17. J. Howard, S. Dighe, Y. Hoskote, S. Van al, D. Finan, G. Ruhl, D. Jenkins, H. Wilson, N. Borkar, G. Schrom, F. Paillet, S. Jain, T. Jacob, S. Yada, S. Marella, P. Salihundam, V. Erraguntla, M. Konow, M. Riepen, G. Droege, J. Lindemann, M. Gries, T. Apel, K. Henriss, T. Lund-larsen, S. Steibl, S. Borkar, V. De, R. Van Der Wijngaart, and T. Mattson. A 48-core ia-32 message-passing processor with DVFS in 45nm CMOS. In *ISSCC '10: Proceedings of the International SolidState Circuits Conference*, 2010.
18. F. D. Igual and G. Quintana-Ortí. Solving linear algebra problems on distributed-memory computers using serial codes. FLAME Working Note #48 DICC 2010-07-01, Universidad Jaume I, Depto. de Ingenieria y Ciencia de Computadores, July 2010.
19. C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh. Basic linear algebra subprograms for Fortran usage. *ACM Trans. Math. Soft.*, 5(3):308–323, Sept. 1979.
20. T. G. Mattson, R. F. Van der Wijngaart, M. Riepen, T. Lehnig, P. Brett, W. Haas, P. Kennedy, J. Howard, S. Van al, N. Borkar, G. Ruhl, and S. Dighe. The 48-core SCC processor: The programmer's view. In *SC'10: Proceedings of the 2010 ACM/IEEE Conference on Supercomputing*, New Orleans, LA, USA, 2010.
21. J. Poulson, B. Marker, J. R. Hammond, N. A. Romero, and R. van de Geijn. Elemental: A new framework for distributed memory dense matrix computations. *ACM Transactions on Mathematical Software*. submitted.
22. G. Quintana-Orti, E. S. Quintana-Orti, R. A. van de Geijn, F. G. Van Zee, and E. Chan. Programming matrix algorithms-by-blocks for thread-level parallelism. *ACM Transactions on Mathematical Software*, 36(3), 2009.
23. P. Sanders and J. L. Träff. The hierarchical factor algorithm for all-to-all communication (research note). In *Proceedings of the 8th International Euro-Par Conference on Parallel Processing*, Euro-Par '02, pages 799–804, London, UK, 2002. Springer-Verlag.
24. R. Schreiber. Scalability of sparse direct solvers. *Graph Theory and Sparse Matrix Computations*, 56, 1992.
25. G. Stewart. Communication and matrix computations on large message passing systems. *Parallel Computing*, 16:27–40, 1990.
26. R. A. van de Geijn. *Using PLAPACK: Parallel Linear Algebra Package*. The MIT Press, 1997.
27. R. A. van de Geijn and E. S. Quintana-Ortí. *The Science of Programming Matrix Computations*. www.lulu.com, 2008.
28. R. Van der Wijngaart, T. G. Mattson, and W. Haas. Light-weight communications on intel's single-chip cloud computer processor. *ACM Operating Systems Review*, 2011. in press.
29. F. G. Van Zee. *libflame: The Complete Reference*. www.lulu.com, 2009.