

Using SIMD Registers and Instructions to Enable Instruction-Level Parallelism in Sorting Algorithms

Timothy Furtak
furtak@cs.ualberta.ca

José Nelson Amaral
amaral@cs.ualberta.ca

Robert Niewiadomski
niewiado@cs.ualberta.ca

Department of Computing Science
University of Alberta, Edmonton, AB, Canada

ABSTRACT

Most contemporary processors offer some version of Single Instruction Multiple Data (SIMD) machinery — vector registers and instructions to manipulate data stored in such registers. The central idea of this paper is to use these SIMD resources to improve the performance of the tail of recursive sorting algorithms. When the number of elements to be sorted reaches a set threshold, data is loaded into the vector registers, manipulated in-register, and the result stored back to memory. Three implementations of sorting with two different SIMD machineries — x86-64's SSE2 and G5's AltiVec — demonstrate that this idea delivers significant speed improvements. The improvements provided are orthogonal to the gains obtained through empirical search for a suitable sorting algorithm [11]. When integrated with the Dynamically Tuned Sorting Library (DTSL) this new code generation strategy reduces the time spent by DTSL up to 22% for moderately-sized arrays, with greater relative reductions for small arrays. Wall-clock performance of *d*-heaps is improved by up to 39% using a similar technique.

Categories and Subject Descriptors

C.1.2 [Processor Architectures]: Multiple Data Stream Architectures (Multiprocessors)—*Single-instruction-stream, multiple-data-stream processors (SIMD)*

General Terms

Algorithms, Performance

Keywords

Quicksort, Sorting, Sorting Networks, SIMD, Instruction-Level Parallelism, Vectorization.

1. INTRODUCTION

This paper addresses the automatic generation of efficient code to sort short sequences of values. The idea is that

an ahead-of-time optimizer searches for fast code for several sequence lengths and machine configurations. Then the compiler can simply instantiate such code when generating an optimized library. While algorithm-specific optimizations and empirical search have long been used both for scientific computation and for large parallel machines [4, 5, 19, 21], only recently these techniques were applied to integer-intensive, symbolic, computation. Li *et al.* developed the Dynamically Tuned Sorting Library that adapts to the characteristics of the input to be sorted [11]. The main contribution of this paper is the insight that the resources implemented in contemporary processors to enable SIMD computations can be put to good use to improve the performance of sorting short sequences. As demonstrated in this work the effective use of these SIMD resources improves performance through the reduction of memory references and increase in instruction level parallelism.

The initial inspiration for this work was the need for fast sorting of short sequences in the implementation of graphics rendering in interactive video-game applications. In such applications it is often necessary to decide, for each pixel of the image, what is the order of the elements that should be displayed [2]. Even though Z-buffer pixel-ordering computations are typically handled by a specialized Graphics Processing Unit (GPU), there are plenty of similar ordering computations that are done by the Central Processing Unit (CPU) in computer games. For instance, sorting is used to characterize the intensity of the various light sources that illuminate a character. Moreover, contemporary video-game application have at their disposal a rich supply of SIMD registers and instructions. For example, the PowerPC-based XBox 360 hardware features 128 AltiVec registers on each of its three cores along with an expanded set of AltiVec instructions. In addition to interactive video-game applications, sorting of short sequences is also present in particle-physics simulation applications.

Thus, using SIMD registers and instructions to sort small sequences is natural. Once a solution was created, applying it to the sequences that must be sorted at the tail-end of standard recursive sorting algorithms was the next logical step. The experimental evaluation of the new vector-register-based sorting algorithms presented in this paper use commodity processors (x86-64 and G5) and extensions to the DTSL library because these machines and algorithms are more readily available and exploitable than proprietary video-game hardware and software. The algorithms presented are effective for sorting short sequences of floating-point or integer values (keys), and pairs comprised of a key

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPAA'07, June 9–11, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-667-7/07/0006 ...\$5.00.

and a memory address, *i.e.* key-pointer pairs, as well as computing the index of a minimum (maximum) element.

Three new SIMD-based algorithms use the concept of sorting networks that are effective to sort small sets of numbers. Section 2 describes: (1) the operation of standard sorting networks; (2) how the SIMD vectors can be used to implement sorting networks; and (3) how a code generator can instantiate optimized vector code for sorting networks operating in sequences of any length. The main contributions in this paper are:

- three algorithms that use the SIMD machinery of contemporary processors for efficient in-register sorting of short sequences and their integration into an optimized general-purpose sorting library;
- a method to use iterative-deepening search to find fast instruction sequences to move data within the SIMD registers;
- a method to compute the minimum element in an array, with applications to d -heaps;
- and an extensive experimental study in three different processors that demonstrate up to 22% improvement in the performance of DTSL for moderately-sized array, and up to 39% in d -heaps. This study also indicates that the elimination of loads, stores, branches, and branch mispredictions correlates well with the improved performance.

Section 3 describes two algorithms that combine a first-pass sorting in the SIMD registers with a second-pass sorting in memory. Section 4 describes an algorithm that sorts shorter sequences completely within the SIMD registers, thus eliminating branch instructions altogether. Section 5 describes how to extend these key-sorting algorithms to sort key-pointer pairs, and Section 6 uses similar techniques to speed up *heapify-down* operations in d -heaps. The experimental evaluation is presented in Section 7.

2. SORTING NETWORKS

The inputs to an in-place comparator, $COMP(a, b)$, are two storage units — memory locations, registers, or vector-register elements — a and b , each containing a numerical input. After the comparator executes, the lower numerical value is stored in a and the higher numerical value is stored in b . Knuth describes a comparator network as a device that applies a fixed sequence of comparator operators to an input vector of a given size [8]. When a comparator network produces a sorted output for any possible input sequence, it is called a sorting network. The size of a sorting network is the total number of comparators in the network. The depth of a sorting network is the length of the critical path in its dependence graph. Therefore the depth provides a bound for the parallel execution of the sorting network, while the size provides a bound for a sequential execution.

An example of a sorting network with size 5 and depth 3 is shown in Fig. 1. The network is depicted as a set of value-carrying vertical rails and comparators. Values flow from top to bottom. A heavy dot at a line crossing indicates that the value at the vertical rail is an input to the comparator represented by the horizontal line. A comparator moves the larger value to the left, and the smaller value to the right.




Figure 1: A 4-element sorting network.

For instance, if the inputs are $a = 7$, $b = 2$, $c = 5$, $d = 9$, then the sorted output at the bottom of the sorting network is $a = 9$, $b = 7$, $c = 5$, and $d = 2$. The value 9 moves from rail d to rail b at $COMP(b, d)$, and then moves from rail b to rail a at $COMP(a, b)$.

Although several algorithms are available to generate code for sorting networks, Batchner’s “odd-even mergesort” algorithm is often chosen for its efficiency [1]. Batchner’s algorithm uses $O(n \log^2 n)$ comparators and has a depth of $O(\log^2 n)$. Sorting networks can be efficiently implemented in processors that provide a *min* and a *max* instruction. Sorting networks implemented with these instructions avoid the performance penalties of branch miss-predictions incurred by traditional branch-based sorting implementations. The experimental results in Section 7 indicate that eliminating branches in the code of sorting networks is a significant win in contemporary processors.

2.1 Supporting Hardware

Consider a machine that has the following *min* and *max* instructions:

$$\min(a, b) = \begin{cases} a : a \leq b \\ b : \text{otherwise} \end{cases}, \max(a, b) = \begin{cases} a : a \geq b \\ b : \text{otherwise} \end{cases}$$

The comparator required by a sorting network is easily constructed using these two operations, a copy instruction, and a temporary variable. For instance, such instructions are available in the x86-64 architectures supporting the SSE2 *min* and *max* operations that return the minimum (maximum) packed single-precision floating-point values [6].¹

The extension of sorting networks to operate on vector instructions requires the definition of vectorized *min* and *max* instructions.² For input vectors A and B , $|A| = |B| = n$, let $C = \min(A, B)$ be the element-wise minimum vector, such that $C_i = \min(A_i, B_i)$, $1 \leq i \leq n$. The vectorized *max* instruction is defined similarly. The width of a (vectorized) sorting network refers to the number of vectors being sorted. Given an ordered list of vectors X^1, X^2, \dots, X^n , a stream of data is formed by selecting the i^{th} element from each vector in order, thus the i^{th} stream is $X_i^1, X_i^2, \dots, X_i^n$.

For instance, the x86-64 architecture has 16 XMM vector registers, and each register can hold 4 floating-point values. Therefore, sorting the values in n XMM registers using a sorting network produces 4 sorted streams of data of length n . Up to 15 XMM registers can be used, *i.e.* $1 \leq n < 16$, because one register must be reserved as temporary storage for the swap of values in the comparator.

This compare-and-swap machinery offers several advantages to sort a small set of values that fits within the SIMD

¹SSE stands for Streaming SIMD Extensions. SSE2 improves upon the original SSE.

²These vector instructions are called a SIMD extension.

registers: (1) its operation is unconditional and data independent; (2) it is inherently branch-free, and thus free of branch-prediction performance penalties; (3) it increases the bandwidth of sorting by enabling the SIMD instruction-level parallelism; and (4) each compare-and-swap requires the execution of only 3 instructions.

A code generator must be able to generate code to sort sequences of any length in a machine with $n + 1$ SIMD registers. The solution is to define size-optimal sorting networks that use $1, 2, \dots, n$ registers. The optimal code for the implementation of each of these sorting networks is pre-generated and stored in a small codebase available to the code generator for deployment. Once data has been loaded into the SIMD registers the code generator instantiates the code to perform the comparator operations specified by the sorting network, and integrates the resulting streams.

3. STREAM-BASED TWO-PASS SORTING

The first two SIMD-based sorting algorithms discussed in this paper operate in two phases. In the first phase the SIMD registers and instructions are used to generate a partially-sorted output. In the second phase a standard sorting algorithm — insertion sort and mergesort are investigated in this paper — finishes the sorting. The choice of algorithm for the second phase dictates the best data organization for the first one.

For the first phase, consider the use of the SIMD sorting machinery described in Section 2 for the task of sorting a sequence of $k * n$ values using n SIMD registers, each register capable of storing k values. Each group of k values is loaded from memory into a separate SIMD register. For a moment, assume that the start of the sequence is aligned for such a load operation. The sorting machinery is then applied to produce k sorted streams of length n , and the sorted streams are written back in-place to memory in an interleaved form. The organization of the data in memory for $k = 4$ is shown in Fig. 2. After sorting, $A_1 \leq A_2 \leq \dots \leq A_n$, $B_1 \leq B_2 \leq \dots \leq B_n$, etc.

After this initial sorting the ordering relationship between elements from separate streams, A_i , B_i , C_i , and D_i , is still unknown. Now the output from the vectorized sorting network must undergo an additional sorting pass. Let us examine the use of insertion sort and mergesort to finish sorting this partially sorted output.

If the start of the sequence is not aligned, the technique used in this paper will be to sort the aligned vector blocks that overlap the target region. The extra fringe elements will be saved to a temporary array, replaced by positive/negative infinity as appropriate, and resorted upon completion.




Figure 2: Interleaved sorted streams from n 4-element SIMD registers. The first register contains elements A_1, B_1, C_1 , and D_1 . $A_1 \leq A_2 \leq \dots \leq A_n$, etc.

3.1 Second Pass with Insertion Sort

A standard insertion-sort algorithm may be used to sort the output of the SIMD-based sorting network. Insertion sort delivers the best performance when its input is mostly sorted because the algorithm does not have to move elements

very far. Thus a potential issue with using insertion sort as a second pass is how the data should be loaded into the SIMD vectors in the first phase to produce the most favorable input for insertion sort.

Consider an input sequence of S values, and a machine with $n + 1$ SIMD vectors. Each vector can store up to k values. Let $m = \lceil S/k \rceil$. If $m \leq n$ the entire array can be loaded into the SIMD registers, sorted, and written back in-place. Then a call to insertion sort will finish sorting the entire sequence.

If $m > n$, an in-place algorithm divides the array into subsets small enough to fit in the vector registers, sorts them with a sorting network, and writes each sorted subset back to the same locations.

A naive approach would simply divide the array into $\lceil m/n \rceil$ almost equal-sized blocks. However, if the data is uniformly distributed this partition results in $\lceil m/n \rceil$ similar blocks, one after the other. The problem is that small elements from the last block would have similar values to the small elements from the first block, and would require insertion sort to move many elements to far positions to combine these blocks.

A better approach is to load the blocks into the SIMD registers in a strided fashion. Consider for example $n = 4$ and $m = 12$ which requires three sorting network calls. Instead of the first call acting on elements A^1, A^2, A^3 , and A^4 , it acts on A^1, A^4, A^7 , and A^{10} . The second call acts on elements A^2, A^5, A^8 , and A^{11} , and the third on A^3, A^6, A^9 , and A^{12} . In this way the small values in the array are likely to end up in A^1, A^2 , and A^3 . A stride width greater than one improves insertion sort performance in cases of uniform or mostly-sorted distributions. In this paper, this strided version of the vectorized sorting network followed by an insertion sort pass is called ISort.

3.2 Second Pass with Mergesort

The mergesort algorithm, called MSort, uses a fixed-sized block of temporary storage T that is large enough to hold the entire array A . Because the SIMD-based sorting is applied to small sequences this array will not be large in practice. MSort proceeds as follows.

Compute the number of blocks of data to be sorted, $\lceil m/n \rceil$, and allocate temporary space T . Call the sorting network on each block from A and store the sorted streams to T .

The *Q-MERGE* algorithm described by Wickremesinghe *et al.* [20] based on work by [14] is now used to store the sorted data into A : (1) Build a heap containing the first element in each stream, and associate with each element a pointer to the next element in its stream; (2) Repeatedly extract the minimum element from the heap. During the extraction, replace the removed element with the next element in its stream, and rebuild the heap.

With a small number of streams, sufficient registers may be available to contain the entire heap. Heapify operations are then efficient and the only flow of data to/from memory is to fetch the next item from a stream or to store the next value to A . For heaps that are too large to fit within the available registers, in-memory heap code may be used. Maintenance operations on small heaps may be written using the known register locations of elements, avoiding potentially costly memory accesses and pointer indirections.

MSort uses one merge heap, with the number of inputs being a multiple of v . That is, each heap completely han-

dles the output from one or more vectorized sorting network calls. Further, only heaps which may be contained within the available registers are considered.

Additional optimizations include placing a sentinel value of infinity at the end of each stream to avoid checking if streams are empty [20]. Once the sentinel is loaded into the head it will sink to the bottom. When any sentinel is extracted from the heap the sorting is complete.

Each sorting network call places elements from the same stream a constant distance away from each other. Thus the next element on a stream can be found by adding a constant offset to the address of the current element, which makes the maintenance of the “next element” pointer in the heap straightforward.

4. ONE-PASS VECTOR SORTING

The third SIMD-based sorting algorithm accomplishes the sorting in a single pass. Intuitively this is possible by loading all of the n elements to be sorted into the vector registers, applying the comparators for an n -element (scalar) sorting network, and writing the elements back to memory in-place.

Table 1: SSE2 instructions used in the example of Fig. 4

Instruction	Description
<code>movaps Ra, Rb</code>	copy the contents of Ra to Rb
<code>shufps Ra, Rb, i</code>	copy 2 elements of Ra to the 2 low-order words of Ra, and 2 elements of Rb to the 2 high-order words of Ra. The elements to be copied are specified by <code>i</code> .
<code>movhyps Ra, Rb</code>	copy the 2 high-order words from Rb to the 2 low-order words of Ra.
<code>movlhps Ra, Rb</code>	copy the 2 low-order words from Rb to the 2 high-order words of Ra.

The difficulty with this approach lies in repositioning elements within the vector registers such that the vector comparator operations do not corrupt the values of elements not involved in the comparison. Moreover, simply aligning comparator inputs may be challenging, depending on the fragmentation of free locations within the vector registers.

Since the cost of applying a vector comparator remains the same regardless of the number of “care” values in each input vector, a natural optimization is to execute more than one (scalar) sorting-network comparator at a time. However, the cost of additional data-movement instructions to properly position multiple comparator inputs in each vector register may outweigh the benefit of parallelization. In practice, for the sorting networks considered, it did not appear to be the case that aligning as many elements as possible³ was ever detrimental to the resulting sequence of operations. However the algorithm we present does provide the ability to balance such alignment costs for the target architecture.

4.1 Searching to Aligning Vector Elements

We will first describe the algorithm used for finding a sequence of alignment instructions, and then show how this applies to a small 4-element sorting network.

³With the optimization that the sorting networks corresponding to Batcher’s Merge Exchange are explicitly separated into layers.




Figure 3: An 8-element sorting network produced by Batcher’s Merge Exchange, with breaks indicated between layers.

4.1.1 Algorithm Input

The input to our algorithm is a sequence of comparators corresponding to a sorting network. In our case the sorting networks were produced by Batcher’s Merge Exchange – not to be confused with Batcher’s Bitonic Sort. Merge Exchange has the property of producing an initial sequence of comparators connecting elements that are separated by powers of 2. This allows for executing a large number of parallel comparators at the start without any need for alignment instructions.

The data dependencies in the sorting network define a partial ordering for the execution of the comparisons. The comparators can thus be partitioned into sets in such a way that all the comparators in each set can be executed in parallel. This partition corresponds to the computation of the maximal anti-chains in a data-dependency graph [18].

One natural optimization, considering multiple legal orderings of the comparator sequence, was not implemented due to the combinatorial increase in the search space. While we present no formal approximation bounds, we feel that the resulting suboptimality of the instruction sequences produced is not significant.

One important optimization which reduces both the number of assembly instructions and the time needed to search for a sequence is to insert explicit breaks between levels of the Merge Exchange sorting network. That is, to disallow executing scalar comparators from different levels within one parallel comparator. For this purpose we consider levels to be the results of the innermost loop in Batcher’s Merge Exchange algorithm as described in [8]. An 8-element Merge Exchange network is shown in Fig. 3 with such layer breaks indicated.

The sequence of alignment instructions within a layer is often repeated for subsequent blocks of scalar comparators. Forcing breaks between levels may be thought of as helping to maintain this repeating pattern of element positions within vectors. This repetition is not exploited directly, but it does seem to introduce less “noise” which may propagate when rearranging elements.

4.1.2 Initial State

For convenience we will assume that we have an unbounded number of vector registers. The resulting sequence of assembly instructions may be restricted to a small number of physical vector registers as a post-processing step by “spilling” and loading values to and from memory as appropriate.

We will also assume that the elements are located in a continuous region of memory, are appropriately aligned, and that the number of elements is a multiple of the size of a vector. These restrictions are for simplification only and may be lifted by making small changes to the algorithm.

Note that the process of searching for a sequence of alignment instructions is only concerned with keeping track of the *labels* of the elements contained within the vector registers – we will refer to manipulations of elements only for convenience.

The first step is to load all of the elements from memory into the vector registers. It is natural and convenient to assume a sequential labeling, such that the first memory location is labeled 0 and the last location $n-1$. Given realistic constraints on the capabilities of the vector manipulation instructions, a number of empty vector registers are required as swap space for rearranging elements. In our experiments having 5 empty vector registers in addition to those registers holding the initial values was seen to be sufficient.

4.1.3 Aligning a Set of Comparators

While all of the comparators in the sorting network have not yet been executed, select the next k comparators that do not cross a layer and such that k is no larger than the number of elements in a vector register. The task is then to rearrange elements such that all the “low” elements from the k comparators are in one vector, and all the “high” elements in another⁴, and aligned element-wise with their partner.

Such an alignment is only valid if applying a vector comparator will not erase the last copy of any element. An erasure must necessarily occur when comparing an element with either an empty (garbage) value or another element with an unknown ordering relation. Note that applying a vector comparator will also invalidate copies of compared elements that are located in other registers.

Finding a sequence of assembly instructions to accomplish this alignment is performed using a standard iterative-deepening search. The legal actions in a state are all vector assembly instructions which do not completely eliminate an element from the set of vector registers.

Due to feasibility concerns, each iterative-deepening search is divided into two phases: moving the low half of each comparator into one vector, and then moving the high half into alignment. If the maximum search depth in any one phase reaches 3, then that task is further subdivided into moving the first 2 elements into a vector, the next 2 elements into another, and finally combining them.

Even with these incremental stages, due to the massive branching factor a naive implementation of this search would take a significant amount of time for even moderately large networks. Our implementation makes use of several admissible heuristics to prune portions of the search space.

To address the tradeoff between the cost of executing a vector comparator and the cost of alignment instructions, the above search is repeated for smaller values of k , and the final cost becomes a combination of the number of alignment instructions and a penalty for including fewer scalar comparators than is possible.

Intuitively this attempts to select the sequence with the best ratio of number of alignments produced versus instruc-

⁴The partitioning of low and high elements may be dropped if relabelling is performed when applying the vector comparator, based on whether the scalar comparator is inverted.

tions required, with an additional bias towards producing more alignments since more alignments will reduce the total number of comparison steps.

When all appropriate values of k have been searched, the choice of how many comparators to include is made greedily and is not revisited. The vector comparator is then applied and the search continues using the remaining comparators.

4.1.4 Writing Values Back to Memory

After the final comparator has been applied the elements are sorted but are not located within the vector registers in an order in which they can be written back to memory. A similar iterative-deepening search now finds an instruction sequence to obtain the correct alignment.

4.2 Example Search

The sorting network shown in Fig. 1 will be used to illustrate the sequence of events in the alignment algorithm for single-pass in-register sorting. This network has four elements and requires the execution of five comparison instructions. An in-register sorting instance of this network using the x86-64 SSE(2) SIMD machinery is shown in Fig. 4. The instructions used in this instance are described in Table 1.⁵

The sorting network of Fig. 1 produces the following partitions: $P_1 = \{COMP(a, c), COMP(b, d)\}$; $P_2 = \{COMP(a, b), COMP(c, d)\}$; and $P_3 = \{COMP(b, c)\}$.

First the elements of **XMM0** are assigned the four elements to be sorted (a , b , c , and d). Then a low-cost sequence of vector instructions is searched for to align a with c and b with d . Here this may be done with a single **movhlps** instruction in step 1. This allows for executing the $COMP(a, c)$ and $COMP(b, d)$ comparators in parallel (step 2)⁶. After this comparison the value stored in element b is smaller than the value stored in element d , and the value stored in element a is smaller than the value stored in element c .

In Fig. 4 a blank square represents a vector element that contains an unknown value that is not relevant to the sorting process. For instance, after the comparison in step 2 the values that were in elements b and a in the low-order words of **XMM0** may have moved. As they are not part of the sorting process they are now represented by blank squares. If the inputs to the sorting network are $a = 7$, $b = 2$, $c = 5$, and $d = 9$, this comparison would leave the highest-order words of **XMM0** and **XMM1** intact and would swap the contents of the second highest-order words. It may also swap the values in the two low-order words of these registers, but the contents of those words are irrelevant.

Now the two comparators in partition P_2 are candidates for the next vector alignment. The initial state for this search is the position of the elements in the vectors at the end of step 2. In the example in Fig. 4 a sequence of two instructions, **movhlps** and **shufps**, is selected to align elements d with c and b with a . Thus both comparators of P_2 can be executed in parallel in step 5.

A penultimate search is performed to execute the last comparator, resulting in steps 6 and 7, at which point the

⁵Other SSE2 instructions frequently used for data movement but not included in this example are: **pshufd**, **unpckhps**, and **unpcklps**.

⁶For SSE2, a comparator between the contents of two registers **Ra** and **Rb** requires a temporary register **T** and the execution of three instructions: **movaps T, Ra**; **minps Ra, Rb**; and **maxps Rb, T**.

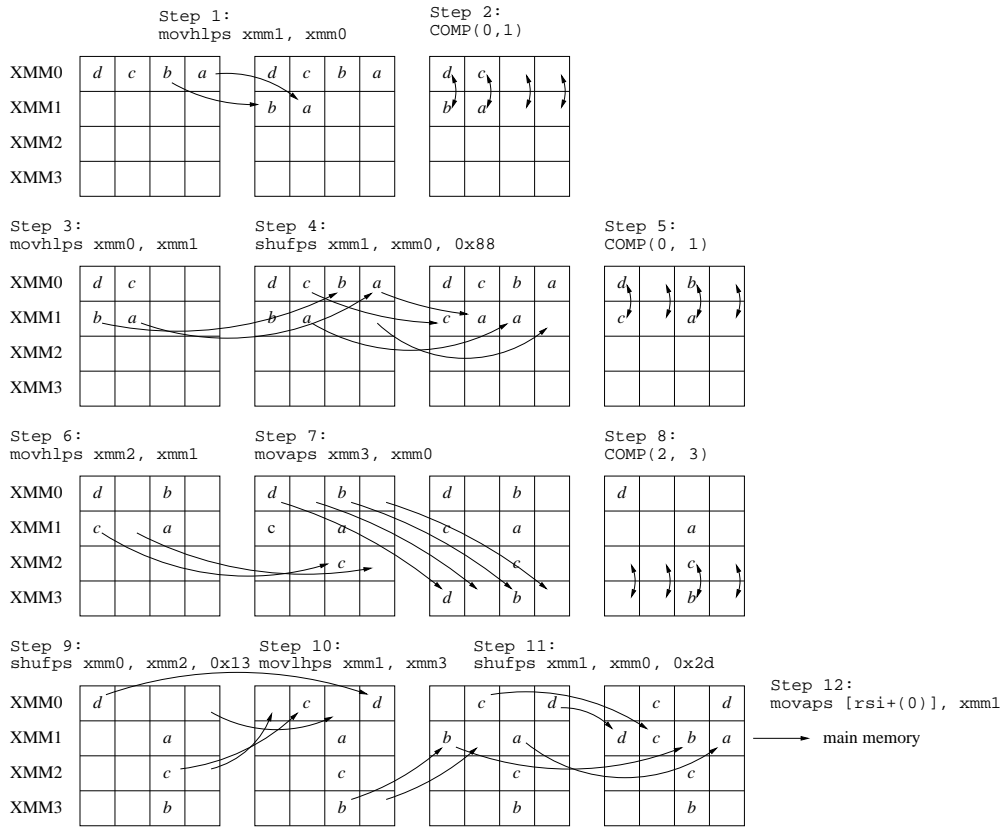


Figure 4: Instruction sequence to apply an in-register 4-element sorting network in an x86-64 architecture. The associated sorting network is shown in Fig. 1.

element values are sorted. Finally, the elements must be properly positioned within one register (in this case XMM1) before the sorted sequence can be written back to memory with a `movaps` instruction.

The vectorization of a sorting network only needs to be done once for each sorting network and for each architecture’s set of vector instructions. Thus all the searches described above should be performed once and offline. The resulting schedule can then be used whenever a sequence of the corresponding size needs to be sorted.

5. SORTING KEY-POINTER PAIRS

So far this paper addresses the problem of sorting an array of floating-point values. A more general problem is that of sorting an array of data structures. Consider the case where each structure has a well-defined floating-point key value. Efficient algorithms sort an array of key-pointer pairs to avoid moving large data structures. This section describes an extension of the vectorized sorting networks to handle key-pointer pairs with floating-point keys and a byte sequence representing the pointer.

The solution to the key-pointer sorting problem consists of storing the keys and the pointers into separate SIMD vectors. If keys and pointers appear interleaved in memory then they must be “swizzled” when loaded into the SIMD vectors and this swizzling must be reversed when storing the sorted result to memory. With the keys and pointers in separate vectors, the standard sorting network solution is

implemented for the keys, while the pointers move in synchrony with the key movements.

This is accomplished by using a bitmask to apply the “swap” operations only to selected elements in the pointer vector. Specifically, those elements which correspond to changes in the key vector after applying the key comparator. The construction of this bitmask is supported in architectures that support SIMD operations. For instance, this may be done in a straightforward manner using the AltiVec `vsel` instruction, while x86-64 architectures must make use of a sequence of boolean operations to mask and combine registers as shown in Fig. 5.

```
asm("pshufd  xmm15, xmm1, 0xE4"); // xmm15 := copy of key_a
asm("minps   xmm1,  xmm2");      // key_a' := min(key_a, key_b)
asm("maxps   xmm2,  xmm15");     // key_b' := max(key_b, key_a)

asm("cmpsps  xmm15, xmm1, 4");   // xmm15 := key_a' != key_a
asm("pshufd  xmm14, xmm3, 0xE4"); // xmm14 := copy of ptr_a
asm("xorps   xmm14, xmm4");      // q := ptr_a XOR ptr_b
asm("andps   xmm15, xmm14");    // q := q AND bitmask
asm("xorps   xmm3,  xmm15");    // ptr_a := ptr_a XOR q
asm("xorps   xmm4,  xmm15");    // ptr_b := ptr_b XOR q
```

Figure 5: Key-pointer comparator using SSE2 assembly instructions. Vector registers xmm1 and xmm2 hold keys, registers xmm3 and xmm4 hold the respective pointers. Registers xmm14 and xmm15 are used as temporary storage.

6. VECTORIZING D-HEAPS

d -heaps are a straightforward generalization of binary heaps where each internal node has d children instead of 2. Increasing the value of d results in a shallower tree at the expense of requiring *delete-min* operations to perform more work when searching for the child node with minimum key value. For concreteness assume *min*-heaps.

Assume an implicit heap layout, with all elements stored in a contiguous array. The root node is located at index 0, and the n th child of a node at index i is located at index $i * d + n$, with $1 \leq n \leq d$. The parent of any node may be similarly computed by dividing its index-1 by d . In [9, 10] LaMarca and Ladner investigate the performance of traditional implicit heaps and how they are affected by data caches. They suggest increasing the branching factor d as well as the data alignment techniques described here and used in our implementation.

We present here a method for increasing d -heap performance by using SIMD vector instructions to quickly compute the index of the child with minimum key value. This computation is used within *heapify-down* operations.

This method is similar to the one used for sorting key-pointer pairs in that it relies on the synchronous movement of values within a second set of registers. In this situation the values moving in synchrony are the *indices* of each child node (specifically the offset from the first child, such that the values range from 0 to $d - 1$).

For simplicity, assume that d is a multiple of k , the number of elements in a SIMD vector. This assumption also aligns a node’s children on both cache-line and SIMD vector boundaries. This alignment requires that the root node be located at the end of a cache-line such that its first child is at the beginning of the next cache-line.

If the nodes in the heap are key-pointer pairs, rather than just keys, loading into a SIMD vector may require additional swizzle instructions to interleave the keys from 2 separate vector loads. Only the key values are required; the associated pointer data may be discarded.

When a block of keys is loaded into a SIMD vector, the index offsets for those keys are loaded into another vector from a constant and static array containing values $0, 1, \dots, d - 1$. The synchronous movement of the index offsets is implemented in the same manner as the movement of the pointer values in Section 5. The loading and movement of these offsets is omitted from the algorithm description for brevity.

The algorithm proceeds as follows: (1) load the first k keys into one SIMD vector, call this register A ; (2) while unread keys remain, read the next k keys into a SIMD vector B and set $A := \min(A, B)$; (3) finally, repeatedly compare one half of the values in A against the other half until only one element remains; (4) return the index of this element. If the node being examined does not have d children (this may only occur at last internal heap node) then the vectorized search is replaced by a straightforward linear scan.

7. EXPERIMENTAL EVALUATION

7.1 Sorting Algorithms

The three versions of vectorized sorting described in this paper were evaluated by integrating them as the low-level algorithms for DTSL’s quicksort. The main findings of this experimental evaluation are:

- Significant reductions in execution time are possible for sorting on the Pentium 4, with lesser reductions on the G5 and Core 2 Duo, depending on array size.
- The integration of SIMD-based sorting algorithms to sort sequences smaller than a fixed threshold improves the performance of DTSL when sorting 51200-element arrays of floating-point key-pointer pairs by up to 22%.
- This performance improvement is due not only to a reduction in the number of loads, stores, and branch instructions, but also to a significant decrease in the number of branch mispredictions.

7.1.1 Integrating Algorithms into DTSL

Table 2: Algorithms studied

Algorithm	Description
MSort $X - Y$	MSort algorithm with X streams applied at Y threshold.
ISort $X - Y$	ISort algorithm with X streams applied at Y threshold.
RSort - Y	One-pass register sort applied at Y threshold.
DTSL - Y	Original DTSL quicksort with SN applied at Y threshold.
Ins - Y	Standard insertion sort applied at Y threshold.




Figure 6: Reduction of branch mispredictions on a 64-bit 3.40 GHz Pentium 4.

The SIMD-based algorithms presented in this paper were integrated in the quicksort implementation of DTSL. The DTSL’s quicksort is not recursive. Instead it maintains an in-function stack of current partitions. When the number of elements to be sorted drops below a threshold, DTSL switches to a low-level sorting algorithm. The version of quicksort that produces the best, or close the best, performance when sorting elements in DTSL uses a scalar sorting network SN as the low-level algorithm [11]. The single-element comparators in this sorting network are written in the C language and use branch instructions to conditionally perform element interchanges. The default threshold to switch to this low-level algorithm is sixteen elements. This version of DTSL’s quicksort is the baseline for the comparative performance study in this paper. Table 2 lists the algorithms used in this performance evaluation. The standard




Figure 7: Quicksort cycle counts relative to DTSL on a 64-bit 3.40 GHz Pentium 4.

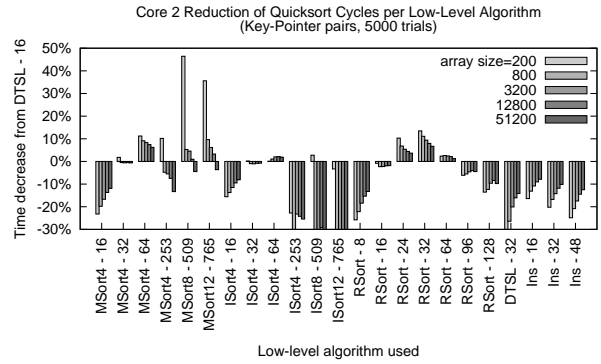


Figure 9: Quicksort cycle counts relative to DTSL on a 3.2 GHz Core 2 Duo E6400.




Figure 8: Quicksort wall-clock times relative to DTSL on a 2.7 GHz Power Mac G5.

insertion-sort algorithm, *Ins - Y*, is included to provide a familiar comparison point.

7.1.2 Wall-Clock Execution Time

Experiments were performed on a 64-bit 3.4 GHz Pentium 4, an IBM 2.7 GHz PowerPC G5, and a 3.2 GHz Core 2 Duo E6400. Figs. 7, 8, 9 show the relative wall-clock execution times for the sorting of a vector of key-pointer pairs in relation to the DTSL baseline. Each bar represents the average runtime over 5000 trials on uniformly distributed keys relative to the DTSL baseline. The large thresholds for MSort, ISort, and RSort, extending beyond what can concurrently fit within the physical vector registers, are the result of the register spilling mentioned in Sec. 4.

Time reductions for the Pentium 4 are quite strong for a range of array sizes, with the greatest reduction of 58% for 200 elements, where all are immediately sorted by *MSort4 - 253*. *RSort - 96* becomes the better alternative for larger arrays, with a time reduction of 22% for 51200 elements.

Large time reductions on the Core 2 Duo and the G5 are limited to small array sizes. For 200 elements *MSort8 - 509* has a respective time reduction of 43% and 33%. For the largest array *RSort - 32* achieves only a 7% and 4% relative improvement on these architectures.

As seen in Fig. 6 the number of branch mispredictions for each algorithm tends to decrease as the threshold becomes larger, reflecting the reduced number (or lack) of branch instructions involved. However larger *RSort* thresholds require

additional functions, much more so than for the two-pass algorithms. These functions grow proportional to the size of their respective sorting networks and include alignment operations. The size of some of the generated object files spans to several megabytes. The first-pass sorting instructions for MSort and ISort do not require nearly as much space.

7.1.3 Low-Level Algorithm Timing

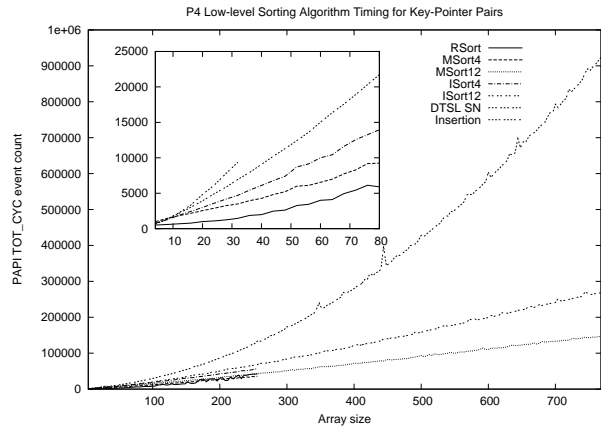


Figure 10: Low-level algorithm cycle counts on a 64-bit 3.40 GHz Pentium 4.

Fig. 10 shows the number of clock cycles, obtained through the PAPI library, required by each algorithm as the number of elements to be sorted varies to the maximum (as implemented) for each algorithm. Each point in the graph is the average over 10000 trials with uniformly distributed keys. This graph shows that *RSort* is significantly superior to both the SN branch-intensive and standard insertion sort, and confirms that *MSort* is also an excellent choice for the sorting of short sequences.

The performance of *RSort* on the G5 and Core 2 is roughly the same as that of the Pentium 4 results shown in Fig. 10 for sequences smaller than 32 elements.

A detailed study of other performance counters showed a correlation between reduction in the number of branches, loads, and stores executed and the relative performance of the algorithms.

7.2 D-Heaps

The performance of d -heaps was investigated by comparing highly optimized versions with different branching factors against SIMD variants where vector instructions were used during *heapify-down* operations. The main findings are a significant reduction in cycle count for larger heaps, when comparing the best SIMD d -heap against the best non-SIMD d -heap.




Figure 12: Ratio of the best SIMD heap times relative to best non-SIMD heap times from Fig. 11.

All source code was written in C++ and was compiled using gcc 3.4.6, 4.0.0, and 4.1.1 on the Pentium 4, G5, and Core 2 Duo respectively, with full optimizations and loop unrolling. The branching factor d was known at compile time. The heap itself was aligned in memory such that the root node’s children began on a cache-line boundary. A consequence is that all children are aligned for SIMD vector accesses. The binary heap had further optimized index computations.

As with the previous experiments, heap elements are key-pointer pairs. Heap have sizes which are powers of 2, from 2^4 to 2^{26} , and are initialized by inserting n elements, where n is the maximum size of the heap. Keys for initial elements are drawn uniformly from $0, \dots, n - 1$.

10,000,000 iterations based on the Hold model as described in [7] were then performed. Each iteration consists of a call to *delete-min* followed by *insert-element*. The key of the new element is equal to the key of element last removed plus a value drawn uniformly from $0, \dots, n - 1$.

As seen in Fig. 11, when the heap size becomes 2^{18} there is a crossover between values of d in the performance of traditional heaps. For small heaps $d = 2$ performs better, while $d = 8$ or $d = 16$ performs better for larger heaps, resulting from better locality of each node’s children as well as decreased heap depth.

All graphs in Fig. 11 show only the best or near-best values of d for clarity. Fig. 12 shows the ratio of execution times between the best SIMD heap at each size versus the best traditional heap. For the Pentium 4, G5, and Core 2 Duo, the SIMD heaps have an average reduction in cycles of 31%, 18%, and 15% respectively, with the largest reductions occurring at the 2^{18} crossover point for the Pentium 4 and G5, and at 2^{20} for the Core 2.

8. RELATED WORK

The implementation of sorting in large-scale vector machines has been extensively studied. Siegel produced one of the earliest descriptions of how to implement Batcher’s sorting network, also known as *bitonic sorting*, in SIMD ma-

chines [17]. Bitton *et al.* provides an extensive description of such implementations [3]. The new contribution of this paper is to demonstrate how the well-known sorting networks can be implemented in the SIMD machinery of contemporary processors and to indicate that code generators can instance such implementations to improve the performance of recursive sorting algorithms and heaps.

The idea of making better use of register resources within the processor to reduce the number of load of stores, in our case to put the SIMD resources to good use in sorting, is also explored by Arge *et al.* [20]. Their idea of forming cache-load-sized runs with quicksort is similar to our idea of switching to SIMD-register-based sorting at an appropriate threshold. The contrast is that we are also benefiting from the SIMD machinery which allows more parallelism in the execution and the elimination of branches while they use the general-purpose registers and the storage available at a cache line.

Recently compilers have been used more often to improve the code generation for SIMD machinery in contemporary processors. Ren *et al.*’s approach of using an optimization algorithm to improve the data permutations is more general than our specific iterative-deepening search [15]. Nuzman *et al.* describes a compiler framework to generate vectorized code for interleaved data [13].

The relationship between the SIMD-register-based sorting algorithms presented in this paper and the development of DTSL is an orthogonal improvement to a library generator [11]. Li *et al.* focused on the dynamic identification of the best sorting algorithm for a given input sequence [12]. They selected an efficient algorithm for the tail of their recursive method. This paper offers a better solution for the sorting of sequences that are small enough to benefit from the use of the SIMD machinery. Similarly, we provide a faster mechanism for selecting a minimum (maximum) child in the implicit d -heaps studied by LaMarca and Ladner [9, 10].

Our SIMD-register-based sorting could also improve partition based sorting methods. For instance, Shen and Ding use an adaptive partitioning scheme to attempt to evenly partition data into chunks smaller than cache size and then use quicksort or insertion sort to finish sorting each bucket [16]. This paper offers a better solution for the sorting of sequences that are small enough to benefit from the use of the SIMD machinery.

9. CONCLUSIONS

This paper proposes the use of the SIMD machinery provided in modern processors to improve the performance of recursion tails. The idea is that whenever the number of elements to be processed fits within the SIMD registers available in the processor, these values should be loaded once into the SIMD registers and then an efficient SIMD execution should be used. While the feasibility of this idea was demonstrated with the integration of a more efficient algorithm for sorting short sequences into DTSL, the idea should be generally applicable to recursive computation.

Once efficient low-level SIMD algorithms are crafted, they can be generated into a solution database to be instantiated by code generators into optimized libraries. Alternatively, if a suitable identification algorithm is created, the compiler should be able to integrate these solutions directly into general programs.

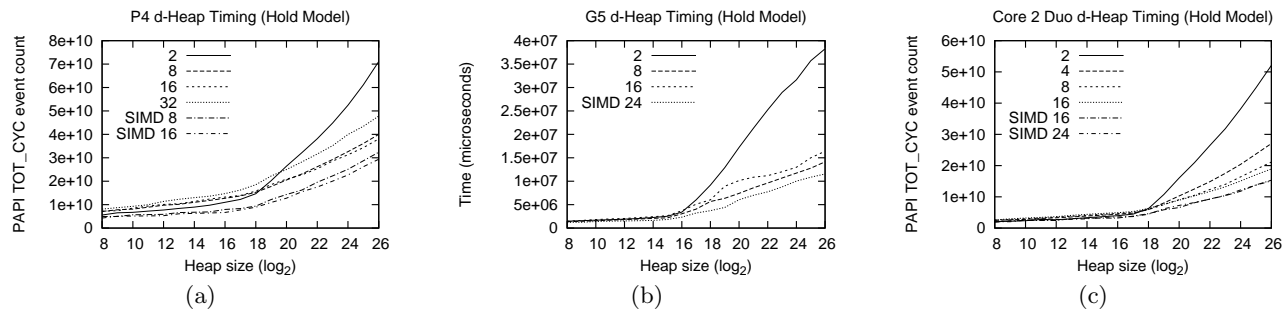


Figure 11: Cycle count / wall-clock time for different heap sizes and values of d on a: (a) 64-bit 3.40 GHz Pentium 4; (b) 2.7 GHz Power Mac G5; (c) 3.20 GHz Core 2 Duo E6400. 10^7 insertions and deletions.

Acknowledgments

The experimental evaluation of these ideas was made possible thanks to David Padua's generous sharing of his group's DTSL code. This research is supported by grants from the Natural Science and Engineering Research Council (NSERC) of Canada, and by IBM Corporation.

10. REFERENCES

- [1] K. E. Batcher. Sorting networks and their applications. In *AFIPS Spring Joint Computing Conference*, pages 307–314, 1968.
- [2] L. Bishop, D. Eberly, T. Whitted, M. Finch, and M. Shantz. Designing a PC game engine. *IEEE Computer Graphics and Applications*, 18(1):46–53, 1998.
- [3] D. Bitton, D. J. DeWitt, D. K. Hsiao, and J. Menon. A taxonomy of parallel sorting. *Computing Surveys*, 16(3):287–318, September 1984.
- [4] J. D. Frens and D. S. Wise. Auto-blocking matrix-multiplication or tracking BLAS3 performance from source code. In *Principles and Practice of Parallel Programming PPOPP*, pages 206–216, Las Vegas, Nevada, 1997.
- [5] M. Frigo. A fast Fourier transform compiler. In *Programming Language Design and Implementation PLDI*, pages 169–180, Atlanta, GA, June 1999.
- [6] Intel. IA-32 Intel®64 and ia-32 architectures software developer's manual volume 1: Basic architecture. <http://www.intel.com/design/processor/manuals/253665.pdf>, 2007.
- [7] Douglas W. Jones. An empirical comparison of priority-queue and event-set implementations. *Commun. ACM*, 29(4):300–311, 1986.
- [8] Donald Ervin Knuth. *The Art of Computer Programming, Vol. 3 - Sorting and Searching*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1973.
- [9] A. LaMarca and R. E. Ladner. The influence of caches on the performance of heaps. *ACM Journal of Experimental Algorithms*, 1:4, 1996.
- [10] A. LaMarca and R. E. Ladner. The influence of caches on the performance of sorting. In *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*, 1997.
- [11] X. Li, M. Garzaran, and D. Padua. A dynamically tuned sorting library. In *Code Generation and Optimization CGO*, pages 111–122, Palo Alto, CA, 2004.
- [12] X. Li, M. J. Garzarán, and D. Padua. Optimizing sorting with genetic algorithms. In *Code Generation and Optimization CGO*, pages 99–110, San Jose, CA, March 2005.
- [13] D. Nuzman, I. Rosen, and A. Zaks. Auto-vectorization of interleaved data for SIMD. In *Programming language design and implementation PLDI*, pages 132–143, 2006.
- [14] A. Ranade, S. Kothari, and R. Udupa. Register efficient mergesorting. In *High Performance Computing — HiPC*, volume 1970 of *LNCS*, pages 96–103. Springer, 2000.
- [15] Gang Ren, Peng Wu, and David Padua. Optimizing data permutations for SIMD devices. In *Programming language design and implementation PLDI*, pages 118–131, 2006.
- [16] Xipeng Shen and Chen Ding. Adaptive data partition for sorting using probability distribution. In *ICPP '04: Proceedings of the 2004 International Conference on Parallel Processing (ICPP'04)*, pages 250–257, Washington, DC, USA, 2004. IEEE Computer Society.
- [17] H. J. Siegel. The universality of various types of SIMD machine interconnection networks. In *Proceedings of the 4th Annual Symposium on Computer Architecture*, pages 23–25, Silver Spring, MD, March 1977. ACM SIGARCH/IEEE-CS.
- [18] S. A. A. Touati. Register saturation in instruction level parallelism. *International Journal of Parallel Programming*, 33(4):393–449, 2005.
- [19] R. Whaley, A. Petitet, and J. Dongarra. Automated empirical optimizations of software and the ATLAS project. *Parallel Computing*, 27(1-2):3–35, 2001.
- [20] R. Wickremesinghe, L. Arge, J. S. Chase, and J. S. Vitter. Efficient sorting using registers and caches. *ACM Journal of Experimental Algorithmics*, 7:9, 2002.
- [21] J. Xiong, J. Johnson, R. Johnson, and D. Padua. SPL: A language and compiler for DSP algorithms. In *Programming Language Design and Implementation PLDI*, pages 298–308, Snowbird, Utah, June 2001.