Exploiting Superword Level Parallelism with Multimedia Instruction Sets

Samuel Larsen and Saman Amarasinghe
MIT Laboratory for Computer Science
Cambridge, MA 02139
{slarsen,saman}@lcs.mit.edu

Abstract

Increasing focus on multimedia applications has prompted the addition of multimedia extensions to most existing general purpose microprocessors. This added functionality comes primarily with the addition of short SIMD instructions. Unfortunately, access to these instructions is limited to in-line assembly and library calls. Generally, it has been assumed that vector compilers provide the most promising means of exploiting multimedia instructions. Although vectorization technology is well understood, it is inherently complex and fragile. In addition, it is incapable of locating SIMD-style parallelism within a basic block.

In this paper we introduce the concept of *Superword Level Parallelism (SLP)*, a novel way of viewing parallelism in multimedia and scientific applications. We believe SLP is fundamentally different from the loop level parallelism exploited by traditional vector processing, and therefore demands a new method of extracting it. We have developed a simple and robust compiler for detecting SLP that targets basic blocks rather than loop nests. As with techniques designed to extract ILP, ours is able to exploit parallelism both across loop iterations and within basic blocks. The result is an algorithm that provides excellent performance in several application domains. In our experiments, dynamic instruction counts were reduced by 46%. Speedups ranged from 1.24 to 6.70.

1 Introduction

The recent shift toward computation-intensive multimedia workloads has resulted in a variety of new multimedia extensions to current microprocessors [6, 10, 16, 18, 20]. Many new designs are targeted specifically at the multimedia domain [3, 7, 11]. This trend is likely to continue as it has been projected that multimedia processing will soon become the main focus of microprocessor design [8].

While different processors vary in the type and number of multimedia instructions offered, at the core of each is a set of short SIMD or superword operations. These instructions operate concurrently on data that are packed in a single register or memory location. In the past, such systems could accommodate only small data types of 8 or 16 bits, making them suitable for a limited set of applications. With the emergence of 128-bit superwords, new architectures are capable of performing four 32-bit operations with a single instruction. By adding floating point support as well, these extensions can now be used to perform more general purpose computation.

It is not surprising that SIMD execution units have appeared in desktop microprocessors. Their simple control, replicated functional units, and absence of heavily-ported register files make them inherently simple and extremely amenable to scaling. As the number of available transistors increases with advances in semiconductor technology, datapaths are likely to grow even larger.

Today, use of multimedia extensions is difficult since application writers are largely restricted to using in-line assembly routines or specialized library calls. The problem is exacerbated by inconsistencies among different instruction sets. One solution to this inconvenience is to employ vectorization techniques that have been used to parallelize scientific code for vector machines [5, 14, 15]. Since a number of multimedia applications are vectorizable, this approach promises good results. However, many important multimedia applications are difficult to vectorize. Complicated loop transformation techniques such as loop fission and scalar expansion are required to parallelize loops that are only partially vectorizable [2, 4, 17]. Consequently, no commercial compiler currently implements this functionality. This paper presents a method for extracting SIMD parallelism beyond vectorizable loops.

We believe that short SIMD operations are well suited to exploit a fundamentally different type of parallelism than the vector parallelism associated with traditional vector and SIMD supercomputers. We denote this parallelism *Superword Level Parallelism (SLP)* since it comes in the form of superwords containing packed data. Vector supercomputers require large amounts of parallelism in order to achieve speedups, whereas SLP can be profitable when parallelism is scarce. From this perspective, we have developed a general algorithm for detecting SLP that targets basic blocks rather than loop nests.

In some respects, superword level parallelism is a restricted form of ILP. ILP techniques have been very successful in the general purpose computing arena, partly because of their ability to find parallelism within basic blocks. In the same way that loop unrolling translates loop level parallelism into ILP, vector parallelism can be transformed into SLP. This realization allows for the parallelization of vector-

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 of the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
PLDI 2000, Vancouver, British Columbia, Canada.
Copyright 2000 ACM 1-58113-199-2/00/0006...$5.00.

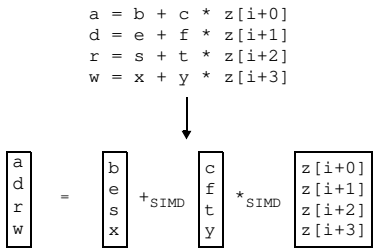

Figure 1: Isomorphic statements that can be packed and executed in parallel.

izable loops using the same basic block analysis. As a result, our algorithm does not require any of the complicated loop transformations typically associated with vectorization. In fact, vector parallelism alone can be uncovered using a simplified version of the SLP compiler algorithm.

The remainder of this paper is organized as follows: Section 2 defines superword level parallelism and compares it to other forms of parallelism. Section 3 describes the compiler algorithm used to extract superword level parallelism. A variation of this algorithm targeting vector parallelism is discussed in Section 4. Section 5 presents results on multimedia and scientific benchmarks. Section 6 discusses architectural features that complement SLP compilation. Section 7 outlines reasons why we believe SLP algorithms will be successful, and Section 8 concludes.

2 Superword Level Parallelism

This section begins by elaborating on the notion of SLP and the means by which it is detected. Terminology is introduced that facilitates the discussion of our algorithms in Sections 3 and 4. We then contrast SLP to other forms of parallelism and discuss their interactions. This helps motivate the need for a new compilation technique.

2.1 Description of Superword Level Parallelism

Superword level parallelism is defined as short SIMD parallelism in which the source and result operands of a SIMD operation are packed in a storage location. Detection is done through a short, simple analysis in which independent isomorphic statements are identified within a basic block. Isomorphic statements are those that contain the same operations in the same order. Such statements can be executed in parallel by a technique we call *statement packing*, an example of which is shown in Figure 1. Here, source operands in corresponding positions have been packed into registers and the addition and multiplication operators have been replaced by their SIMD counterparts. Since the result of the computation is also packed, unpacking may be required depending on how the data are used in later computations. The performance benefit of statement packing is determined by the speedup gained from parallelization minus the cost of packing and unpacking.

Depending on what operations an architecture provides to facilitate general packing and unpacking, this technique can actually result in a performance degradation if packing and unpacking costs are high relative to ALU operations. One of the main objectives of our SLP detection technique is to minimize packing and unpacking by locating cases in which packed data produced as a result of one computation can be used directly as a source in another computation.

```
for (i=0; i<16; i++) {
  localdiff = ref[i] - curr[i];
  diff += abs(localdiff);
}
```

(a) Original loop.

```
for (i=0; i<16; i++) {
  T[i] = ref[i] - curr[i];
}

for (i=0; i<16; i++) {
  diff += abs(T[i]);
}
```

(b) After scalar expansion and loop fission.

```
for (i=0; i<16; i+=4) {
  localdiff = ref[i+0] - curr[i+0];
  diff += abs(localdiff);

  localdiff = ref[i+1] - curr[i+1];
  diff += abs(localdiff);

  localdiff = ref[i+2] - curr[i+2];
  diff += abs(localdiff);

  localdiff = ref[i+3] - curr[i+3];
  diff += abs(localdiff);
}
```

(c) Superword level parallelism exposed after unrolling.

```
for (i=0; i<16; i+=4) {
  localdiff0 = ref[i+0] - curr[i+0];
  localdiff1 = ref[i+1] - curr[i+1];
  localdiff2 = ref[i+2] - curr[i+2];
  localdiff3 = ref[i+3] - curr[i+3];

  diff += abs(localdiff0);
  diff += abs(localdiff1);
  diff += abs(localdiff2);
  diff += abs(localdiff3);
}
```

(d) Packable statements grouped together after renaming.

Figure 2: A comparison between SLP and vector parallelization techniques.

Packed statements that contain adjacent memory references among corresponding operands are particularly well suited for SLP execution. This is because operands are effectively pre-packed in memory and require no reshuffling within a register. In addition, an address calculation followed by a load or store need only be executed once instead of individually for each element. The combined effect can lead to a significant performance increase. This is not surprising since vector machines have been successful at exploiting the same phenomenon. In our experiments, instructions eliminated from operating on adjacent memory locations had the greatest impact on speedup. For this reason, locating adjacent memory references forms the basis of our algorithm, discussed in Section 3.

2.2 Vector Parallelism

Vector parallelism is a subset of superword level parallelism. Our results in Section 5 show that 20% of dynamic instruction savings on the SPEC95fp benchmark suite are from non-vectorizable code sequences.

To better explain the differences between superword level parallelism and vector parallelism, we present two short examples, shown in Figures 2 and 3. Although the first ex-

```
do {
  dst[0] = (src1[0] + src2[0]) >> 1;
  dst[1] = (src1[1] + src2[1]) >> 1;
  dst[2] = (src1[2] + src2[2]) >> 1;
  dst[3] = (src1[3] + src2[3]) >> 1;

  dst  += 4;
  src1 += 4;
  src2 += 4;
}
while (dst != end);
```

Figure 3: An example of a hand-optimized matrix operation that proves unvectorizable.

ample can be molded into a vectorizable form, we know of no vector compilers that can be used to vectorize the second. Furthermore, the transformations required in the first example are unnecessarily complex and may not work in more complicated circumstances. In general, a vector compiler must employ a repertoire of tools in order to parallelize loops on a case by case basis. By comparison, our method is simple and robust, yet still capable of detecting the available parallelism.

Figure 2(a) presents the inner loop of the motion estimation algorithm used for MPEG encoding. Vectorization is inhibited by the presence of a loop carried dependence and a function call within the loop body. To overcome this, a vector compiler can perform a series of transformations to mold the loop into a vectorizable form. The first is scalar expansion, which allocates a new element in a temporary array for each iteration of the loop [4]. Loop fission is then used to divide the statements into separate loops [12]. The result of these transformations is shown in Figure 2(b). The first loop is vectorizable, but the second must be executed sequentially.

Figure 2(c) shows the loop from the perspective of SLP. After unrolling, the four statements corresponding to the first statement in the original loop can be packed together. The packing process effectively moves packable statements to contiguous positions, as shown in part (d). The code motion is legal because it does not violate any dependences (once scalar renaming is performed). The first four statements in the resulting loop body can be packed and executed in parallel. Their results are then unpacked so they can be used in the sequential computation of the final statements. In the end, this method has the same effect as the transformations used for vector compilation, while only requiring loop unrolling and scalar renaming.

Figure 3 shows a code segment that averages the elements of two 16x16 matrices. As is the case with many multimedia kernels, our example has been hand-optimized for a sequential machine. In order to vectorize this loop, a vector compiler would need to reverse the programmer-applied optimizations. Were such methods available, they would involve constructing a *for* loop, restoring the induction variable, and re-rolling the loop. In contrast, locating SLP within the loop body is simple. Since the optimized code is amenable to SLP analysis, hand-optimization has had no detrimental effects on our ability to detect the available parallelism.

2.3 Loop Level Parallelism

Vector parallelism, exploited by vector computers, is a subset of loop level parallelism. General loop level parallelism is typically exploited by a multiprocessor or MIMD machine. In many cases, parallel loops may not yield performance gains because of fine-grain synchronization or loop-carried communication. It is therefore necessary to find coarse-grain parallel loops when compiling for MIMD machines. Traditionally, a MIMD machine is composed of multiple microprocessors. It is conceivable that loop level parallelism could be exploited orthogonally to superword level parallelism within each processor. Since coarse-grain parallelism is required to get good MIMD performance, extracting SLP should not detract from existing MIMD parallel performance.

2.4 SIMD Parallelism

SIMD parallelism came into prominence with the advent of massively parallel supercomputers such as the Illiac IV [9]. The association of the term "SIMD" with this type of computer is what led us to utilize the term Superword Level Parallelism when discussing short SIMD operations.

SIMD supercomputers were implemented using thousands of small processors that worked synchronously on a single instruction stream. While the cost of massive SIMD parallel execution and near-neighbor communication was low, distribution of data to these processors was expensive. For this reason, automatic SIMD parallelization centered on solving the data distribution problem [1]. In the end, the class of applications for which SIMD compilers were successful was even more restrictive than that of vector and MIMD machines.

2.5 Instruction Level Parallelism

Superword level parallelism is closely related to ILP. In fact, SLP can be viewed as a subset of instruction level parallelism. Most processors that support SLP also support ILP in the form of superscalar execution. Because of their similarities, methods for locating SLP and ILP may extract the same information. Under circumstances where these types of parallelism completely overlap, SLP execution is preferred because it provides a less expensive and more energy efficient solution.

In practice, the majority of ILP is found in the presence of loops. Therefore, unrolling the loop multiple times may provide enough parallelism to satisfy both ILP and SLP processor utilization. In this situation, ILP performance would not noticeably degrade after SLP is extracted from a program.

3 SLP Compiler Algorithm

Our SLP compiler algorithm can be divided into several distinct phases. First, loop unrolling is used to transform vector parallelism into SLP. Alignment analysis then attempts to determine the address alignment of each load and store instruction. This is needed for compiling to architectures that do not support unaligned memory accesses. Next, the intermediate representation is transformed into a low level form and a series of standard compiler optimizations is applied.

The core of our algorithm begins by locating statements with adjacent memory references and packing them into groups of size two. From this initial seed, more groups are discovered based on the active set of packed data. All groups are then merged into larger clusters of a size consistent with the superword datapath width. Finally, a new schedule is

produced for each basic block, where groups of packed statements are replaced with SIMD instructions.

The following subsections describe each of these phases in detail. Figure 4 presents a simple example to highlight the core routines and Figure 5 lists the pseudo code. Both will be referenced throughout this section.

3.1 Loop Unrolling

Loop unrolling is performed early since it is most easily done at a high level. As discussed, it is used to transform vector parallelism into basic blocks with superword level parallelism. In order to ensure full utilization of the superword datapath in the presence of a vectorizable loop, the unroll factor must be customized to the data sizes used within the loop. For example, a vectorizable loop containing 16-bit values should be unrolled 8 times for a 128-bit datapath. Our system currently unrolls loops based on the smallest data type present.

3.2 Alignment Analysis

Alignment analysis determines the alignment of memory accesses with respect to a certain superword datapath width. For architectures that do not support unaligned memory accesses, alignment analysis can greatly improve the performance of our system. Without it, memory accesses are assumed to be unaligned and the proper merging code must be emitted for every wide load and store.

One situation in which merging overhead can be amortized is when a contiguous block of memory is accessed within a loop. In this situation, overhead can be reduced to one additional merge operation per load or store by using data from previous iterations.

Alignment analysis, however, can completely remove this overhead. For FORTRAN sources, a simple interprocedural analysis can determine alignment information in a single pass. This analysis is flow-insensitive, context-insensitive, and visits the call graph in breadth-first order. For C sources, we use an enhanced pointer analysis package developed by Rugina and Rinard [21]. Since this pass also provides location set information, we can consider dependences more carefully when combining packing candidates. A full discussion of alignment analysis is beyond the scope of this paper. A complete description will be given in [13].

Our compilation system is capable of operating both with and without alignment constraints. For simplicity, we describe subsequent phases of the algorithm assuming no architectural support for unaligned accesses. As such, later phases assume alignment information has been annotated to each load and store instruction where possible.

3.3 Pre-optimization

SLP analysis is most useful when performed on a three address representation. This way, the algorithm has full flexibility in choosing which operations to pack. If isomorphic statements are instead matched by the tree structure inherited from the source code, long expressions must be identical in order to parallelize. On the other hand, identifying adjacent memory references is much easier if address calculations maintain their original form. We therefore annotate each load and store instruction with this information before flattening.

After flattening, several standard optimizations are applied to an input program. This ensures that parallelism is not extracted from computation that would otherwise be eliminated. Optimizations include constant propagation, copy propagation, dead code elimination, common subexpression elimination, loop-invariant code motion, and redundant load/store elimination. As a final step, scalar renaming is performed to remove output and anti-dependences since they can inhibit parallelization.

3.4 Identifying Adjacent Memory References

Because of their obvious impact, statements containing adjacent memory references are the first candidates for packing. We therefore begin the core of our analysis by scanning each basic block to find independent pairs of such statements. Adjacency is determined using both alignment information and array analysis.

In general, duplicate memory operations can introduce several different packing possibilities. Dependences will eliminate many of these possibilities and redundant load/store elimination will usually remove the rest. In practice, nearly every memory reference is directly adjacent to at most two other references. These correspond to the references that access memory on either side of the reference in question. When located, the first occurrence of each pair is added to the *PackSet*.

Definition 3.1 *A Pack is an n-tuple,* $\langle s_1, ..., s_n \rangle$*, where* $s_1, ..., s_n$ *are independent isomorphic statements in a basic block.*

Definition 3.2 *A PackSet is a set of Packs.*

In this phase of the algorithm, only groups of two statements are constructed. We refer to these as *pairs* with a *left* and *right* element.

Definition 3.3 *A Pair is a Pack of size two, where the first statement is considered the left element, and the second statement is considered the right element.*

As an intermediate step, statements are allowed to belong to two groups as long as they occupy a *left* position in one of the groups and a *right* position in the other. Enforcing this discipline here allows the *Combination* phase to easily merge groups into larger clusters. These details are discussed in Section 3.6.

Figure 4(a) presents an example sequence of statements. Figure 4(b) shows the results of adjacent memory identification in which two pairs have been added to the *PackSet*. The pseudo code for this phase is shown in Figure 5 as find_adj_refs.

3.5 Extending the PackSet

Once the *PackSet* has been seeded with an initial set of packed statements, more groups can be added. This is done by finding new candidates that can either:

- Produce needed source operands in packed form, or
- Use existing packed data as source operands.

This is accomplished by following def-use and use-def chains of existing *PackSet* entries. If these chains lead to fresh packable statements, a new group is created and added to the *PackSet*. For two statements to be packable, they must meet the following criteria:

- The statements are isomorphic.

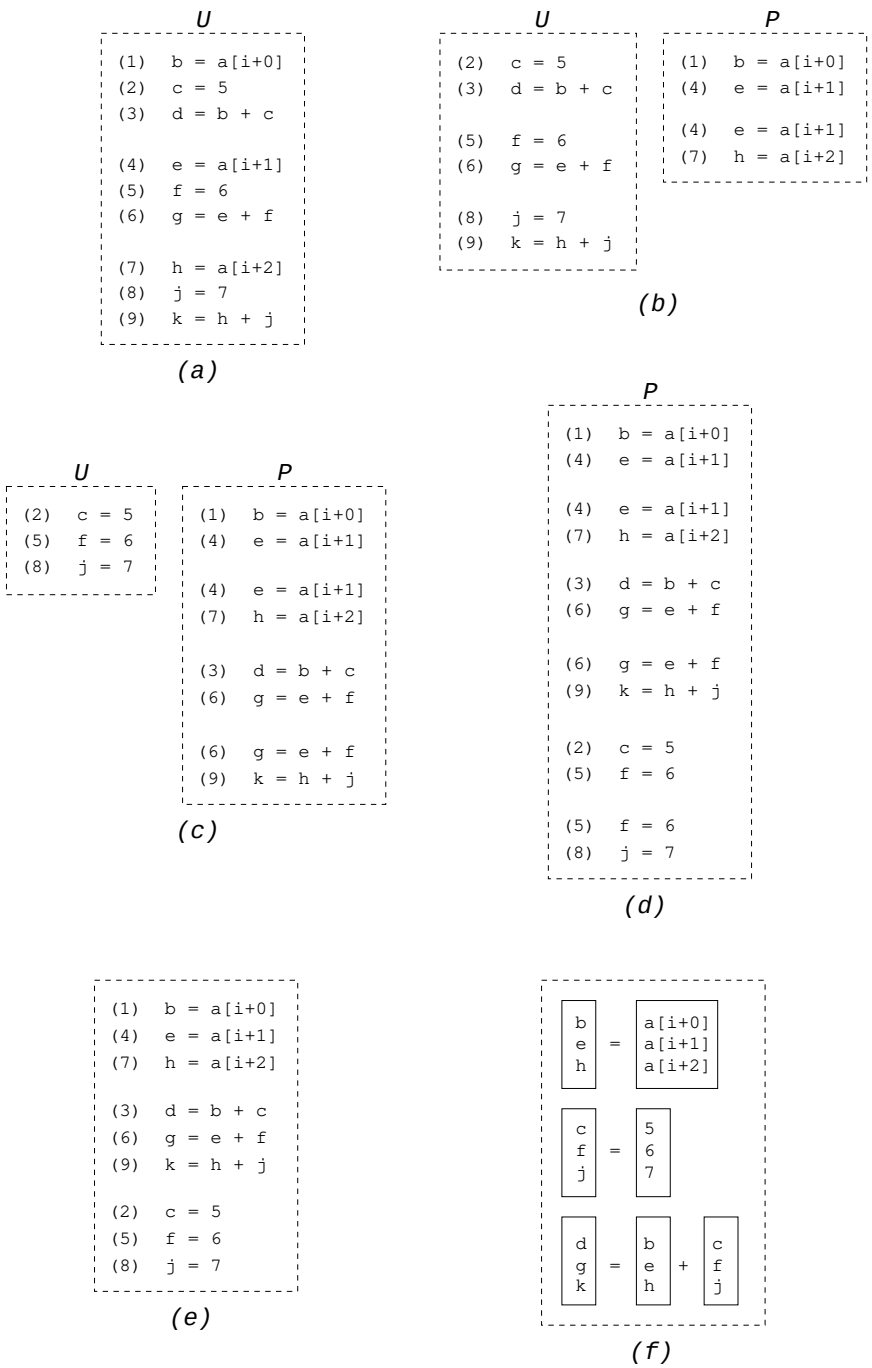

Figure 4: Various phases of SLP analysis. U and P represent the current set of unpacked and packed statements, respectively. (a) Initial sequence of instructions. (b) Statements with adjacent memory references are paired and added to the *PackSet*. (c) The *PackSet* is extended by following def-use chains of existing entries. (d) The *PackSet* is further extended by following use-def chains. (e) *Combination* merges groups containing the same expression. (f) Each group is scheduled and SIMD operations are emitted in their place.

```
SLP_extract: BasicBlock B → BasicBlock
    PackSet P ← ∅
    P ← find_adj_refs(B, P)
    P ← extend_packlist(B, P)
    P ← combine_packs(P)
    return schedule(B, [ ], P)

find_adj_refs: BasicBlock B × PackSet P → PackSet
    foreach Stmt s ∈ B do
      foreach Stmt s' ∈ B where s ≠ s' do
        if has_mem_ref(s) ∧ has_mem_ref(s') then
          if adjacent(s, s') then
            Int align ← get_alignment(s)
            if stmts_can_pack(B, P, s, s', align) then
              P ← P ∪ {⟨s, s'⟩}
    return P

extend_packlist: BasicBlock B × PackSet P → PackSet
    repeat
      PackSet P_prev ← P
      foreach Pack p ∈ P do
        P ← follow_use_defs(B, P, p)
        P ← follow_def_uses(B, P, p)
    until P ≡ P_prev
    return P

combine_packs: PackSet P → PackSet
    repeat
      PackSet P_prev ← P
      foreach Pack p = ⟨s_1, ..., s_n⟩ ∈ P do
        foreach Pack p' = ⟨s'_1, ..., s'_m⟩ ∈ P do
          if s_n ≡ s'_1 then
            P ← P − {p, p'} ∪ {⟨s_1, ..., s_n, s'_2, ..., s'_m⟩}
    until P ≡ P_prev
    return P

schedule: BasicBlock B × BasicBlock B' × PackSet P
        → BasicBlock
    for i ← 0 to |B| do
      if ∃p = ⟨..., s_i, ...⟩ ∈ P then
        if ∀s ∈ p. deps_scheduled(s, B') then
          foreach Stmt s ∈ p do
            B ← B − s
            B' ← B' · s
          return schedule(B, B', P)
        else if deps_scheduled(s_i, B') then
          return schedule(B − s_i, B' · s_i, P)
    if |B| ≠ 0 then
      P ← P − {p} where p = first(B, P)
      return schedule(B, B', P)
    return B'

stmts_can_pack: BasicBlock B × PackSet P ×
          Stmt s × Stmt s' × Int align → Boolean
    if isomorphic(s, s') then
      if independent(s, s') then
        if ∀⟨t, t'⟩ ∈ P.t ≠ s then
          if ∀⟨t, t'⟩ ∈ P.t' ≠ s' then
            Int align_s ← get_alignment(s)
            Int align_s' ← get_alignment(s')
            if align_s ≡ ⊤ ∨ align_s ≡ align then
              if align_s' ≡ ⊤ ∨ align_s' ≡ align+data_size(s') then
                return true
    return false

follow_use_defs: BasicBlock B × PackSet P × Pack p → PackSet
          where p = ⟨s, s'⟩, s = [ x_0 := f(x_1, ..., x_m) ], s' = [ x'_0 := f(x'_1, ..., x'_m) ]
    Int align ← get_alignment(s)
    for j ← 1 to m do
      if ∃t ∈ B.t = [ x_j := ... ] ∧ ∃t' ∈ B.t' = [ x'_j := ... ] then
        if stmts_can_pack(B, P, t, t', align)
          if est_savings (⟨t, t'⟩, P) ≥ 0 then
            P ← P ∪ {⟨t, t'⟩}
            set_alignment(s, s', align)
    return P

follow_def_uses: BasicBlock B × PackSet P × Pack p → PackSet
          where p = ⟨s, s'⟩, s = [ x_0 := f(x_1, ..., x_m) ], s' = [ x'_0 := f(x'_1, ..., x'_m) ]
    Int align ← get_alignment(s)
    Int savings ← −1
    foreach Stmt t ∈ B where t = [ ... := g(..., x_0, ...) ] do
      foreach Stmt t' ∈ B where t ≠ t' = [ ... := h(..., x'_0, ...) ] do
        if stmts_can_pack(B, P, t, t', align) then
          if est_savings(⟨t, t'⟩, P) > savings then
            savings ← est_savings(⟨t, t'⟩, P)
            Stmt u ← t
            Stmt u' ← t'
    if savings ≥ 0 then
      P ← P ∪ {⟨u, u'⟩}
      set_alignment(u, u')
    return P
```

Figure 5: Pseudo code for the SLP extraction algorithm. Only key procedures are listed. Helper functions include: 1) has_mem_ref, which returns true if a statement accesses memory, 2) adjacent, which checks adjacency between two memory references, 3) get_alignment, which retrieves alignment information, 4) set_alignment, which sets alignment information when it is not already set, 5) deps_scheduled, which returns true when, for a given statement, all statements upon which it is dependent have been scheduled, 6) first, which returns the *PackSet* member containing the earliest unscheduled statement, 7) est_savings, which estimates the savings of a potential group, 8) isomorphic, which checks for statement isomorphism, and 9) independent, which returns true when two statements are independent.

- The statements are independent.
- The left statement is not already packed in a *left* position.
- The right statement is not already packed in a *right* position.
- Alignment information is consistent.
- Execution time of the new parallel operation is estimated to be less than the sequential version.

The analysis computes an estimated speedup of each potential SIMD instruction based on a cost model for each instruction added and removed. This includes any packing or unpacking that must be performed in conjunction with the new instruction. If the proper packed operand data already exist in the *PackSet*, then packing cost is set to zero.

As new groups are added to the *PackSet*, alignment information is propagated from existing groups via use-def or def-use chains. Once set, a statement's alignment determines which position it will occupy in the datapath during its computation. For this reason, a statement can have only one alignment. New groups are created only if their alignment requirements are consistent with those already in place.

When a single definition has multiple uses, there is the potential for many different packing possibilities. If this occurs, the cost model is used to estimate the most profitable possibilities based on what is currently packed. These groups are added to the *PackSet* in order of their estimated profitability as long as there are no conflicts with existing *PackSet* entries.

In the example, part (c) shows new groups that are added after following def-use chains of the two existing *PackSet* entries. Part (d) introduces new groups discovered by following use-def chains. The pseudo code for this phase is listed as extend_packset in Figure 5.

3.6 Combination

Once all profitable pairs have been chosen, they can be combined into larger groups. Two groups can be combined when the *left* statement of one is the same as the *right* statement of the other. In fact, groups must be combined in this fashion in order to prevent a statement from appearing in more than one group in the final *PackSet*. This process, provided by the combine_packs routine, checks all groups against one another and repeats until all possible combinations have been made. Figure 4(e) shows the result of our example after combination.

Since the adjacent memory identification phase uses alignment information, it will never create pairs of memory accesses that cross an alignment boundary. All packed statements are aligned based on this initial seed. As a result, the combination phase will never produce a group that spans an alignment boundary. Combined groups are therefore guaranteed to be less than or equal to the superword datapath size.

3.7 Scheduling

Dependence analysis before packing ensures that statements within a group can be executed safely in parallel. However, it may be the case that executing two groups produces a dependence violation. An example of this is shown in Figure 6. Here, dependence edges are drawn between groups

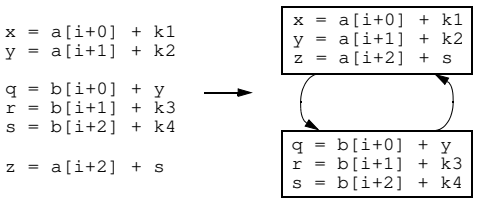

Figure 6: Example of a dependence between groups of packed statements.

if a statement in one group is dependent on a statement in the other. As long as there are no cycles in this dependence graph, all groups can be scheduled such that no violations occur. However, a cycle indicates that the set of chosen groups is invalid and at least one group will need to be eliminated. Although experimental data has shown this case to be extremely rare, care must be taken to ensure correctness.

The scheduling phase begins by scheduling statements based on their order in the original basic block. Each statement is scheduled as soon as all statements on which it is dependent have been scheduled. For groups of packed statements, this property must be satisfied for each statement in the group. If scheduling is ever inhibited by the presence of a cycle, the group containing the earliest unscheduled statement is split apart. Scheduling continues until all statements have been scheduled.

Whenever a group of packed statements is scheduled, a new SIMD operation is emitted instead. If this new operation requires operand packing or reshuffling, the necessary operations are scheduled first. Similarly, if any statements require unpacking of their source data, the required steps are taken. Since our analysis operates at the level of basic blocks, each basic block assumes all data are in an unpacked configuration upon entry to the block. For this reason, all variables that are live on exit are unpacked at the end of the block.

Scheduling is provided by the schedule routine in Figure 5. In the example of Figure 4, the result of scheduling is shown in part (f). At the completion of this phase, a new basic block has been constructed wherever parallelization was successful. These blocks contain SIMD instructions in place of packed isomorphic statements. As we will show in Section 5, the algorithm can be used to achieve speedups on a microprocessor with multimedia extensions.

4 A Simple Vectorizing Compiler

The SLP concepts presented in the previous section lead to an elegant implementation of a vectorizing compiler. Vector parallelism is characterized by the execution of multiple iterations of an instruction using a single vector operation. This same computation can be uncovered with unrolling by limiting packing to unrolled versions of the same statement. With this technique, each statement has only one possible grouping, which means that no searching is required. Instead, every statement can be packed automatically with its siblings if they are found to be independent. The profitability of each group can then be evaluated in the context of the entire set of packed data. Any groups that are deemed unprofitable can be dropped in favor of their sequential counterparts. The pseudo code for this algorithm is shown in Figure 7.

While not as general as the algorithm described in the

```
vector_parallelize: BasicBlock B → BasicBlock
    PackSet P ← ∅
    P ← find_all_packs(B, P)
    P ← eliminate_unprofitable_packs(P)
    return schedule(B, [ ], P)


find_all_packs: BasicBlock B × PackSet P → PackSet
    foreach Stmt s ∈ B do
      if ∀p ∈ P.s ∉ p then
        Pack p ← [s]
        foreach Stmt s' ∈ B where s' ≠ s do
          if stmts_are_packable(s, s') then
            p ← p · s'
        if |p| > 1 then
          P ← P ∪ {p}
    return P


stmts_are_packable: Stmt s × Stmt s' → Boolean
    if same_orig_stmt(s, s') then
      if independent(s, s') then
        return true
    return false


eliminate_unprofitable_packs: PackSet P → PackSet
    repeat
      PackSet P' ← P
      foreach Pack p ∈ P do
        if est_savings(p, P) < 0 then
          P ← P − {p}
    until P ≡ P'
    return P
```

Figure 7: Pseudo code for the vector extraction algorithm. Procedures that are identical to those in Figure 5 are omitted. same_orig_stmt returns true if two statements are unrolled versions of the same original statement.

previous section, this technique shares many of the same desirable properties. First, the analysis itself is extremely simple and robust. Second, partially vectorizable loops can be parallelized without complicated loop transformations. Most importantly, this analysis is able to achieve good results on scientific and multimedia benchmarks.

The drawback to this method is that it may not be applicable to long vector architectures. Since the unroll factor must be consistent with the vector size, unrolling may produce basic blocks that overwhelm the analysis and the code generator. As such, this method is mainly applicable to architectures with short vectors.

In Section 5, we will provide data that compare this approach to the algorithm described in Section 3.

5 Results

This section presents potential performance gains for SLP compiler techniques and substantiates them using a Motorola MPC7400 microprocessor with the AltiVec instruction set. All results were gathered using the compiler algorithms described in Sections 3 and 4. Both were implemented within the SUIF compiler infrastructure [23].

5.1 Benchmarks

We measure the success of our SLP algorithm on both scientific and multimedia applications. For scientific codes, we use the SPEC95fp benchmark suite. Our multimedia benchmarks are provided by the kernels listed in Table 1.

Name	Description	Datatype
FIR	Finite impulse response filter	32-bit float
IIR	Infinite impulse response filter	32-bit float
VMM	Vector-matrix multiply	32-bit float
MMM	Matrix-matrix multiply	32-bit float
YUV	RGB to YUV conversion	16-bit integer

Table 1: Multimedia kernels used to evaluate the effectiveness of SLP analysis.

5.2 SLP Availability

To evaluate the availability of superword level parallelism in our benchmarks, we calculated the percentage of dynamic instructions eliminated from a sequential program after parallelization. All instructions were counted equally, including SIMD operations. When packing was required, we assumed that *n-1* instructions were required to pack n values into a single SIMD register. These values were also used for unpacking costs.

Measurements were obtained by instrumenting source code with counters in order to determine the number of times each basic block was executed. These numbers were then multiplied by the number of static SUIF instructions in each basic block. Results for both sets of benchmarks are listed in Table 2 and illustrated in Figure 8. The performance of each benchmark is shown for a variety of hypothetical datapath widths. It is assumed that each datapath can accommodate SIMD versions of any standard data type. For example, a datapath of 512 bits can perform eight 64-bit floating point operations in parallel. To uncover the maximum amount of superword level parallelism available, we compiled each benchmark without alignment constraints. This allowed for a maximum degree of freedom when making packing decisions.

For the multimedia benchmarks, YUV greatly outperforms the other kernels. This is because it operates on 16-bit values and is entirely vectorizable. The remaining kernels are partially vectorizable and still exhibit large performance gains.

For the SPEC95fp benchmark suite, some of the appli-

Benchmark	128 bits	256 bits	512 bits	1024 bits
swim	61.59%	64.45%	73.44%	77.17%
tomcatv	40.91%	61.28%	69.50%	73.85%
mgrid	43.49%	55.13%	60.51%	61.52%
su2cor	33.99%	48.73%	56.06%	59.63%
wave5	26.69%	37.25%	41.97%	43.87%
apsi	24.19%	29.93%	31.32%	29.85%
hydro2d	18.53%	26.17%	28.88%	30.80%
turb3d	21.16%	24.76%	21.55%	15.13%
applu	15.54%	22.56%	10.29%	0.01%
fpppp	4.22%	8.14%	8.27%	8.27%
FIR	38.72%	45.37%	48.56%	49.84%
IIR	51.83%	60.59%	64.77%	66.45%
VMM	36.92%	43.37%	46.63%	51.90%
MMM	61.75%	73.63%	79.76%	82.86%
YUV	87.21%	93.59%	96.79%	98.36%

Table 2: Percentage of dynamic instructions eliminated by SLP analysis for a variety of hypothetical datapath widths.

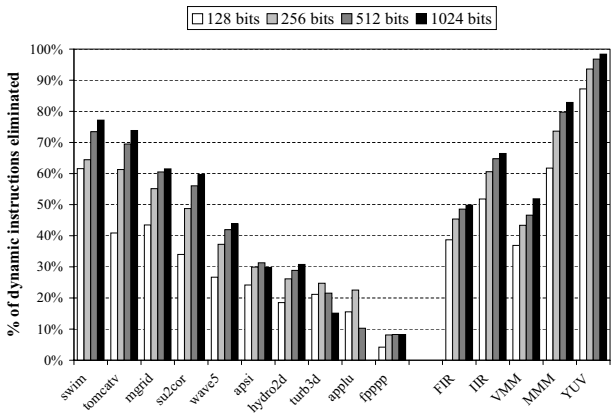

Figure 8: Percentage of dynamic instructions eliminated by SLP analysis for a variety of hypothetical datapath widths.

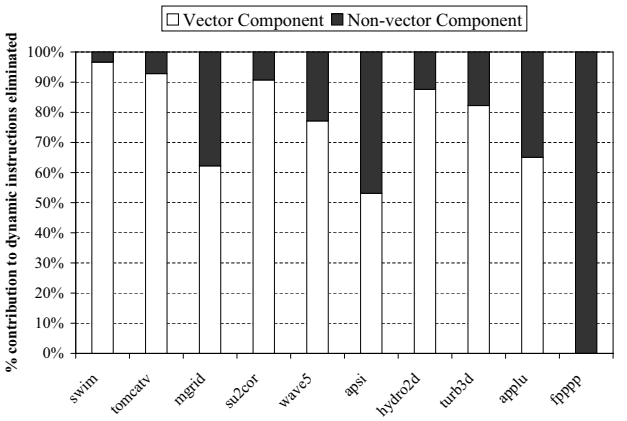

Figure 10: Contribution of vectorizable and non-vectorizable code sequences in total SLP savings for the SPEC95fp benchmark suite.

cations exhibit a performance degradation as the datapath width is increased. This is due to the large unroll factor required to fill a wide datapath. If the dynamic iteration counts for these loops are smaller than the unroll factor, the unrolled loop is never executed. For turb3d and applu, the optimal unroll factor is four. A 256-bit datapath is therefore sufficient since it can accommodate four 64-bit operations. In fpppp, the most time-intensive loop is already unrolled by a factor of three. A 192-bit datapath can support the available parallelism in this situation.

In Figure 9 and Table 3 we compare the SLP algorithm to the vectorization technique described in Section 4. For the multimedia benchmarks, both methods perform identically. However, there are many cases in the scientific applications for which the SLP algorithm is able to find additional packing opportunities. In Figure 10, we show the available vector parallelism as a subset of the available superword level parallelism.

5.3 SLP Performance

To test the performance of our SLP algorithm in a real environment, we targeted our compilation system to the AltiVec [19] instruction set. Of the popular multimedia extensions available in commercial microprocessors, we believe AltiVec best matches the compilation technique described in this paper. AltiVec defines 128-bit floating point and integer SIMD operations and provides a complementary set of 32 general purpose registers. It also defines load and store instructions capable of moving a full 128 bits of data.

Our compiler automatically generates C code with AltiVec macros inserted where parallelization is successful. We then use an extended gcc compiler to generate machine code. This compiler was provided by Motorola and supports the AltiVec ABI (application binary interface). Due to the experimental nature of the AltiVec compiler extensions, it was necessary to compile all benchmarks without optimization. Base measurements were made by compiling the unparallelized version for execution on the MPC7400 superscalar unit. In both cases, the same set of SUIF optimizations and the same gcc backend were used. Since AltiVec does not support unaligned memory accesses, all benchmarks were compiled with alignment constraints in place [13].

Table 4 and Figure 11 present performance comparisons on a 450MHz G4 PowerMac workstation. Most of the SPEC95fp benchmarks require double precision floating point support to operate correctly. Since this is not

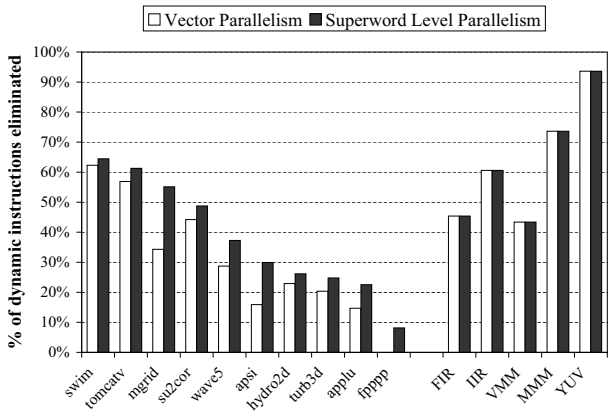

Figure 9: Percentage of dynamic instructions eliminated with SLP parallelization and with vector parallelization on a 256-bit datapath.

Benchmark	SLP	Vector
swim	64.45%	62.29%
tomcatv	61.28%	56.87%
mgrid	55.13%	34.29%
su2cor	48.73%	44.20%
wave5	37.25%	28.73%
apsi	29.93%	15.89%
hydro2d	26.17%	22.91%
turb3d	24.76%	20.35%
applu	22.56%	14.67%
fpppp	8.14%	0.00%
FIR	45.37%	45.37%
IIR	60.59%	60.59%
VMM	43.37%	43.37%
MMM	73.63%	73.63%
YUV	93.59%	93.59%

Table 3: Percentage of dynamic instructions eliminated with SLP parallelization and with vector parallelization on a 256-bit datapath.

Benchmark	Speedup
swim	1.24
tomcatv	1.57
FIR	1.26
IIR	1.41
VMM	1.70
MMM	1.79
YUV	6.70

Table 4: Speedup on an MPC7400 processor using SLP compilation.

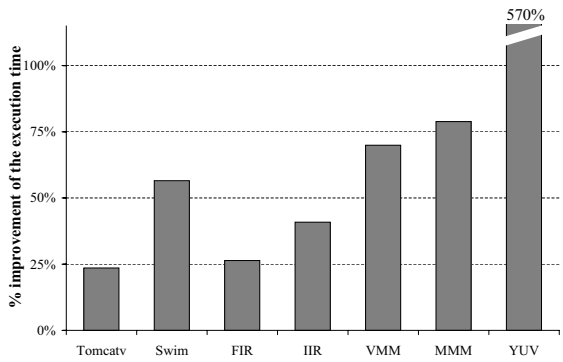

Figure 11: Percentage improvement of execution time on an MPC7400 processor using SLP compilation.

supported by AltiVec, we were unable to compile vectorized versions for all but two of the benchmarks. `swim` utilizes single precision floating point operations, and the SPEC92fp version of `tomcatv` provides a result similar to the 64-bit version.

Our compiler currently assumes that all packed operations are executed on the AltiVec unit and all sequential operations are performed on the superscalar unit. Operations to pack and unpack data are therefore required to go through memory since AltiVec provides no instructions to move data between register files. Despite this high cost, our compiler is still able to exploit superword level parallelism and provide speedups.

6 Architectural Support for SLP

The compiler algorithm presented in Section 3 was inspired by the multimedia extensions in modern processors. However, several limitations make it difficult to fully realize the potential provided by SLP analysis. We list some of these limitations below:

- Many multimedia instructions are designed for a specific high-level operation. For example, HP's MAX-2 extensions offer matrix transform instructions [16] and SUN's VIS extensions include instructions to compute pixel distances [18]. The complex CISC-like semantics of these instructions make automatic code generation difficult.
- SLP hardware is typically viewed as a multimedia engine alone and is not designed for general purpose computation. Floating point capabilities, for example, have only recently been added to some architectures. Furthermore, even the most advanced multimedia extensions lack certain fundamental operations such as 32-bit integer multiplication and division [19].
- In current architectures, data sets are usually considered to belong exclusively to either multimedia or superscalar hardware. This design philosophy is portrayed in the lack of inter register file move operations in the AltiVec instruction set. If SLP compilation techniques can show a need for a better coupling between these two units, future architectures may provide the necessary support.
- Most current multimedia instruction sets are designed with the assumption that data are always stored in the proper packed configuration. As a result, data packing and unpacking instructions are generally not well supported. This important operation is useful to our system. With better support, SLP performance can be further increased.
- Although our system is capable of compiling for machines that do not support unaligned memory accesses, the algorithm is potentially more effective without this constraint. Architectures supplying efficient unaligned load and store instructions might improve the performance of SLP analysis.

The first three points discuss simple processor modifications that we hope will be incorporated into future multimedia instruction sets as they mature. The last two points address difficult issues. Solving them in either hardware or software is not trivial. More research is required to determine the best approach.

7 Keys to General Acceptance of SLP

Many of the techniques developed by the academic compiler community are not accepted in mainstream computing. A good example is the work on loop level parallelization that has continued for over three decades. However, in a very short period of time, ILP compilers have become universal. We believe the following characteristics are critical to the general acceptance of a compiler optimization:

- **Robustness:** If simple source code modifications drastically alter program performance, success becomes dependent upon the user's understanding of compiler intricacies. For example, techniques to uncover loop level parallelism are prone to wide fluctuations in performance. A change in one statement of the loop body may result in a vector compiler's sequentialization of the entire loop. In the case of ILP and SLP, failure to parallelize a few statements will not significantly impact aggregate performance. This makes methods for their extraction much more robust.
- **Scalability:** Compiler techniques must be able to handle large programs if they are to gain acceptance for real applications. Some analyses required by loop optimizations do not scale well to large code sizes because of dependence on global program analysis. Although global analysis can improve the effectiveness of ILP and SLP, it is not required. Therefore, complexity grows linearly with program size. This results in smooth scaling to larger applications.

- **Simplicity:** Complex compiler transformations are more prone to bugs than simple analyses. Problems are likely to appear only under very specific conditions, making them difficult to detect. Many time-critical projects are compiled without optimizations in order to avoid possible compiler errors. Coarse-grain parallelization and vectorization require involved analyses that are more likely to exhibit this behavior. However, most ILP techniques, as well as the SLP techniques presented in Section 3, are extremely simple to understand, implement and validate. In addition, it is often the case that simplicity leads to faster compilation.

- **Portability:** Optimizations that are dependent on particular features of a source language or programming style will not become universal. Techniques for extracting loop level parallelism are limited because they only apply to programs written with loops and arrays. Alternatively, ILP and SLP techniques are applied at the level of basic blocks, making them less dependent on source code characteristics.

- **Effectiveness:** No compiler technique will be used if it does not substantially improve program performance. In Section 5, we showed that our algorithm for detecting SLP can provide remarkable performance gains.

We believe SLP compiler techniques have the potential to become universally accepted as viable and effective methods of extracting SIMD parallelism. As a result, we expect future architectures to place increasing importance on SLP operations.

8 Conclusion

In this paper we introduced superword level parallelism, the notion of viewing parallelism from the perspective of partitioned operations on packed superwords. We showed that SLP can be exploited with a simple and robust compiler implementation that exhibits speedups ranging from 1.24 to 6.70 on a set of scientific and multimedia benchmarks.

We also showed that SLP concepts lead to an elegant implementation of a vectorizing compiler. By comparing the performance of this compiler to the more general SLP algorithm, we demonstrated that vector parallelism is a subset of superword level parallelism.

Our current compiler implementation is still in its infancy. While successful, we believe its effectiveness can be improved. By extending the SLP analysis beyond basic blocks, more packing opportunities could be found. Furthermore, SLP could offer a form of predication, in which unfilled slots of a wide operation could be filled with speculative computation. If data are invalidated due to control flow, they could simply be discarded.

Recent research has shown that compiler analysis can significantly reduce the size of data types needed to store program variables [22]. Incorporating this analysis into our own has the potential of drastically improving performance by increasing the number of operands that can be packed and executed in parallel.

Today, most desktop processors are equipped with multimedia extensions. Nonuniformities in the different instruction sets, exacerbated by a lack of compiler support, has left these extensions underutilized. We have shown that SLP compilation is not only possible, but also applicable to a wider class of application domains. As such, we believe SLP compilation techniques have the potential to become an integral part of general purpose computing in the near future.

9 Acknowledgments

We thank Krste Asanović for his input on vector processing and the CAG group members who provided feedback during the writing of this paper, particularly Michael Taylor, Derek Bruening, Mike Zhang, Darko Marinov, Matt Frank and Mark Stephenson. Radu Rugina contributed his pointer analysis package as well as his own time to implement new extensions. Manas Mandal, Kalpesh Gala, Brian Grayson and James Yang at Motorola provided access to much needed development tools and technical expertise. We also thank the anonymous reviewers for their constructive comments. Finally, we thank Matt Deeds for his help in the completion of this paper.

This research was funded in part by NSF grant EIA9810173 and DARPA grant DBT63-96-C-0036.

References

[1] E. Albert, K. Knobe, J. Lukas, and G. Steele, Jr. Compiling Fortran 8x array features for the Connection Machine computer system. In *Proceedings of the ACM SIGPLAN Symposium on Parallel Programming: Experience with Applications, Languages, and Systems (PPEALS)*, New Haven, CT, July 1988.

[2] J. R. Allen and K. Kennedy. PFC: A Program to Convert Fortran to Parallel Form. In K. Hwang, editor, *Supercomputers: Design and Applications*, pages 186–203. IEEE Computer Society Press, Silver Spring, MD, 1984.

[3] Krste Asanović, James Beck, Bertrand Irissou, Brian E. D. Kingsbury, Nelson Morgan, and John Wawrzynek. The T0 Vector Microprocessor. In *Proceedings of Hot Chips VII*, August 1995.

[4] D. Callahan and P. Havlak. Scalar expansion in PFC: Modifications for Parallelization. Supercomputer Software Newsletter 5, Dept. of Computer Science, Rice University, October 1986.

[5] Derek J. DeVries. A Vectorizing SUIF Compiler: Implementation and Performance. Master's thesis, University of Toronto, June 1997.

[6] Keith Diefendorff. Pentium III = Pentium II + SSE. *Microprocessor Report*, 13(3):1,6–11, March 1999.

[7] Keith Diefendorff. Sony's Emotionally Charged Chip. *Microprocessor Report*, 13(5):1,6–11, April 1999.

[8] Keith Diefendorff and Pradeep K. Dubey. How Multimedia Workloads Will Change Processor Design. *IEEE Computer*, 30(9):43–45, September 1997.

[9] G. H. Barnes, R. Brown, M. Kato, D. J. Kuck, D. L. Slotnick, and R. A. Stokes. The Illiac IV Computer. *IEEE Transactions on Computers*, C(17):746–757, August 1968.

[10] Linley Gwennap. AltiVec Vectorizes PowerPC. *Microprocessor Report*, 12(6):1,6–9, May 1998.

[11] Craig Hansen. MicroUnity's MediaProcessor Architecture. *IEEE Micro*, 16(4):34–41, Aug 1996.

[12] D.J. Kuck, R.H. Kuhn, D. Padua, B. Leasure, and M. Wolfe. Dependence Graphs and Compiler Optimizations. In *Proceedings of the 8th ACM Symposium on Priciples of Programming Languages*, pages 207–218, Williamsburg, VA, Jan 1981.

[13] Samuel Larsen, Radu Rugina, and Saman Amarasinghe. Alignment Analysis. Technical Report LCS-TM-605, Massachusetts Institute of Technology, June 2000.

[14] Corina G. Lee and Derek J. DeVries. Initial Results on the Performance and Cost of Vector Microprocessors. In *Proceedings of the 30th Annual International Symposium on MicroArchitecutre*, pages 171–182, Research Triangle Park, USA, December 1997.

[15] Corina G. Lee and Mark G. Stoodley. Simple Vector Microprocessors for Multimedia Applications. In *Proceedings of the 31st Annual International Symposium on MicroArchitecutre*, pages 25–36, Dallas, TX, December 1998.

[16] Ruby Lee. Subword Parallelism with MAX-2. *IEEE Micro*, 16(4):51–59, Aug 1996.

[17] Glenn Luecke and Waqar Haque. Evaluation of Fortran Vector Compilers and Preprocessors. *Software—Practice and Experience*, 21(9), September 1991.

[18] Marc Tremblay and Michael O'Connor and Venkatesh Narayanan and Liang He. VIS Speeds New Media Processing. *IEEE Micro*, 16(4):10–20, Aug 1996.

[19] Motorola. *AltiVec Technology Programming Environments Manual*, November 1998.

[20] Alex Peleg and Uri Weiser. MMX Technology Extension to Intel Architecture. *IEEE Micro*, 16(4):42–50, Aug 1996.

[21] Radu Rugina and Martin Rinard. Pointer Analysis for Multithreaded Programs. In *Proceedings of the SIGPLAN '99 Conference on Programming Language Design and Implementation*, Atlanta, GA, May 1999.

[22] Mark Stephenson, Jonathon Babb, and Saman Amarasinghe. Bitwidth Analysis with Application to Silicon Compilation. In *Proceedings of the SIGPLAN '00 Conference on Programming Language Design and Implementation*, Vancouver, BC, June 2000.

[23] R. P. Wilson, R. S. French, C. S. Wilson, S. P. Amarasinghe, J. M. Anderson, S. W. K. Tjiang, S.-W. Liao, C.-W. Tseng, M. W. Hall, M. S. Lam, and J. L. Hennessy. SUIF: An Infrastructure for Research on Parallelizing and Optimizing Compilers. *ACM SIGPLAN Notices*, 29(12):31–37, December 1994.