

Filter Fusion

Todd A. Proebsting*
University of Arizona

Scott A. Watterson
University of Arizona

1 Introduction

Filters are a common data-manipulation abstraction in networking, operating systems, and simulation software. Filters read data from a single source and write data to a single destination. In filter applications, data flows from a source to a sink through intermediate filters. Logically, filters are separate, modular entities. Modular implementations unfortunately suffer a substantial performance penalty relative to integrated implementations. Where performance matters most, systems programmers will sacrifice the modular design for the greater speed of an integrated design.

We present a new compiler optimization, Filter Fusion, that eliminates the overhead of a modular design of independent filters. Our algorithm automates the integration of arbitrary, independently designed filters. FFC, our Filter Fusion compiler, composes filters and produces code that is as efficient as hand-integrated code. The optimized code can achieve up to a two-fold improvement over independent filters.

Network protocol layers are often filters. Typically, each protocol layer performs some data manipulation by traversing the message from beginning to end. Programmers have traditionally merged these filters by hand to produce efficient code. Integrating filters allows data to be read once, manipulated many times, and then stored once, thus avoiding loads and stores for each filter's manipulations. Excessive memory ac-

cesses cripple the performance of network code. Filter Fusion eliminates unnecessary memory accesses.

Manually integrating filters is a time-consuming, error-prone process. In addition, hand-integrated programs are difficult to maintain and modify because small changes in a single filter can result in global changes in the integrated program. FFC automates the integration process and therefore eliminates this concern. Furthermore, automatic integration enables the maintenance of a library of useful filters (protocol layers) that can be composed freely to develop specialized protocols. Each library component is maintained separately, and yet integration and optimization is automatic. The programmer designs and optimizes in a modular fashion, without sacrificing performance in the final composition.

While Filter Fusion is well suited for systems software applications, no assumptions about its problem domain are made. FFC places few restrictions on the filters it integrates. It handles arbitrary control flow and data manipulations within each filter.

2 Related Work

2.1 Network Programming

FFC is part of the compiler suite of the Scout project [MMO⁺95]. Scout aims to deliver high-performance systems software—especially communications-oriented operating systems. The Scout compilers do non-traditional optimizations, like Filter Fusion, to increase software performance and to liberate the programmer from tedious, error-prone tasks [OPM94].

Network applications often require many simple manipulations of each network packet. These manipulations form the protocol stack. Redundant memory access can dominate the processing time for these applications. A technique called Integrated Layer Processing (ILP) optimizes these data ma-

Address: Department of Computer Science, University of Arizona, Tucson, AZ 85721. Internet: {todd,saw}@cs.arizona.edu.

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

POPL '96, St. Petersburg FLA USA
© 1996 ACM 0-89791-769-3/95/01...\$3.50

manipulations [CT90]. ILP, a generalization of loop jamming or loop fusion, does increase performance [CT90, CJRS89, DAPP93].

Clark and Tennenhouse report dramatic performance improvements from ILP [CT90]. Based on their results, they argue for *less* modular programming—when efficiency is critical and sequential data manipulations are too costly, the programmer must abandon abstraction and merge protocols. By automating ILP, Filter Fusion allows the programmer to retain modular design without sacrificing performance.

Abbot partially automated ILP for network applications [Abb93]. His system has two significant drawbacks, however: it cannot handle arbitrary control-flow within a filter, and it assumes the typical network data layout that partitions *header* and *data*. His protocols had three stages: initial, data manipulation, and final. The integrated code performed the initial and final stages serially with only the data manipulation stages truly integrated. Not all protocols (e.g., message re-assembly), and certainly not all filters, fit into this framework. Filter Fusion has no such restrictions.

These prior implementations have proven the efficacy of ILP, but they have not fully generalized or automated the optimization. Thus, a tension exists between modular software design and integrated high-performance implementation. FFC, an implementation of Filter Fusion, provides a solution. While maintaining a clean, intuitive model for protocol construction, it provides both modularity and performance.

2.2 Program Transformation

Many different program transformations that improve performance by removing redundant computations have been previously studied. Partial evaluation, listlessness, and deforestation exploit, in one way or another, the basic unfold/instantiate/fold framework originally proposed by Burstall and Darlington to improve programs [BD77]. While this system automated program transformation rather than program analysis, its framework was revolutionary.

Recently, partial evaluation has been the dominant paradigm for eliminating unnecessary computation from programs [JGS93]. While an oversimplification, partial evaluation strengthens Burstall and Darlington’s work by maintaining significant state information—the state of the static values—at all program points during transformation. Partial evaluation has been applied to functional languages with much more success than to imperative languages.

Filter Fusion is a very limited special-purpose partial evaluation system for producer and consumer functions that alternate their computations like coroutines. Partially evaluating adjacent filters with respect to a given composition results in a fusion of those filters. Restricting the optimization to composing filters simplifies the partial evaluation considerably.

Filter Fusion most strongly resembles Wadler’s *listless* transformer [Wad84]. This functional programming optimization composes functions that create and utilize lists into code that avoids building intermediate lists. Like listlessness, Filter Fusion symbolically executes programs to create a graph representations of the residual, composed program that eliminates intermediate steps (list manipulations for listlessness, and reads/writes for Filter Fusion). Filter Fusion, on the other hand, operates on imperative “filters” that explicitly read from and write to one another. Given a function, F , that maps imperative programs to functional programs, and an imperative filter program, P , it would be interesting to know whether $Listlessness(F(P))$ is isomorphic to $F(FilterFusion(P))$. Whether or not Filter Fusion is isomorphic to listlessness is not clear, but the symbolic execution and residual graphs do indicate a close relationship. We developed Filter Fusion without knowledge of the listless transformer.

Deforestation is related to listlessness and Filter Fusion, because it too eliminates intermediate structures through symbolic execution [Wad90].

3 Filters

A linear composition of filters specifies the path data will follow from source to sink:

$$Source \rightarrow Filter_1 \rightarrow \dots \rightarrow Filter_N \rightarrow Sink$$

In a modular implementation, the source produces *all* of the data before passing it to the first filter. That filter then processes all the data before passing it to the next filter. This continues until the sink ultimately consumes the data. Unfortunately, this implementation requires that each filter read and write data. It is much more efficient to merge these filters to perform all the data manipulations at once.

3.1 Filter Specifications

A filter specification is simply a parameterless procedure extended by three operations: **put**, **get**, and **filter**. A **put** produces data for the next filter, and a **get** retrieves data from the previous filter. (Filter

Fusion will merge filters so that matching **put**'s and **get**'s can be replaced by assignments.) **filter** is a special predicate that guides Filter Fusion. **filter** guards statements that either require more input or may produce more output. **filter** is explained further in section 5.

The first filter of a composition, the *source*, cannot contain any **get**'s. The last filter, the *sink*, cannot contain any **put**'s. Figure 1 contains source and sink filters for simple array reading and writing.

Data manipulation filters exist between the source and the sink. Typical filters may do encryption, compression, checksumming, or data marshaling (e.g., byte swapping). In addition, *glue* filters are useful for combining filters that may require special invariants. For instance, the simple filter for swapping pairs of adjacent bytes, **2ByteSwap**, requires an even number of bytes as input. The **Evener** is a glue filter that always writes an even number of bytes by simply copying its input to its output and conditionally appending a single zero. Thus, the **Evener** typically precedes **2ByteSwap** to ensure proper functioning. Figure 2 gives the specifications for **2ByteSwap** and **Evener**. Lightweight filter design encourages modular design and separation of concerns.

Typical network protocols such as CRC32 checksum and MD-5 encryption are also filters. Other functions we have implemented as filters include Run-length Decoding and Run-length Encoding, simple checksumming, and data marshaling. Filter Fusion allows the programmer to create arbitrarily complex compositions of these independently developed filters; **FFC** will integrate them into a single optimized function.

Efficiency and modularity are advantages of using **FFC**. Without **FFC**, reorganizing a protocol stack requires re-integrating the stack by hand. With **FFC**, reorganizing a stack simply requires changing the individual filters (if necessary) and specifying a new composition.

4 Sample Fusion

Filter Fusion is an optimization based on a symbolic execution of the filters. Filter Fusion integrates two filters—a *producer* and a *consumer*—at a time. The goal is to match the **put**'s of the producer with the **get**'s of the consumer and to replace them with assignments. Using dynamic programming, Filter Fusion follows all possible control flow paths through both filters while tracking the flow of values via the **put**'s and **get**'s. Filter Fusion composes the control-flow graphs of the filters into new, larger graph.

Where necessary, Filter Fusion replicates filter code.

As an example, we will merge the **Evener** and the **2ByteSwap** filters in Figure 2. Figure 3 gives their control-flow graphs. Rectangles denote nodes from **2ByteSwap** throughout this example; ovals denote **Evener** nodes.

The final control-flow graph is composed of nodes from the two original graphs, except that the appropriate **put**'s and **get**'s are replaced with assignments to temporary variables. Basically, the dynamic programming executes each filter symbolically—alternating between the producer and consumer at **put**'s and **get**'s, respectively. For each node that is symbolically executed, a copy of that node is placed into the fused graph. Bookkeeping information maintained at each node of the final graph controls the composition. Each added node is annotated with three pieces of information: the last node executed in the producer, the last node executed in the consumer, and which filter this node came from. This information is a *configuration*. Two nodes are equal if their configurations are identical.

The producer symbolically executes until it reaches a **put** or **end** operation. After reaching a **put** in the producer, execution switches to the consumer, which must execute until it reaches a **get** (or **end**). The **put** that suspended the producer is matched with the consumer's **get** for subsequent replacement by an assignment. This alternating execution continues until all possible execution paths are exhausted.

The **filter** predicate will represent a conditional node in a control flow graph of either the producer or the consumer. The state of a suspended producer determines the value of a consumer's **filter** predicate. If a consumer is executing while the producer is suspended at a **put**, then **filter** evaluates to true; if the producer is suspended at its end, then **filter** evaluates to false. **filter** predicates in the producer remain undetermined.¹

Figure 4 depicts the control flow of the fused filter after the producer has followed all possible paths to **put**'s or **end**'s. Symbolic execution must now switch to the consumer.

When expanding the consumer (**2ByteSwap**), the first node to be executed is a **filter** predicate. Thus, all three paths will add a **filter** node. On the left-most path, the producer had suspended at a **end**, but on the center and right-most paths, the producer suspended at a **put**. Therefore, consumer will continue along the false branch when expanding the left-most

¹This discussion assumes that the producer is driving Filter Fusion. If the consumer were driving Filter Fusion, then the **filter** predicates in the producer would be determined by whether or not the consumer were suspended at a **get**.

<pre> Filter ReadFromArray Decls int i; Code i = 0; while (i < 10000) put input[i]; i++; end-while End-Filter </pre>	<pre> Filter WriteToArray Decls int j; Code while filter get output[j]; j++; end-while End-Filter </pre>
---	--

Figure 1: Source and Sink Filters

<pre> Filter Evener Decls int c, k; Code k = 0; while filter get c; put c; k++; end-while if (k%2) put 0; End-Filter </pre>	<pre> Filter 2ByteSwap Decls int x, y; Code while filter get x; get y; put y; put x; end-while End-Filter </pre>
---	--

Figure 2: Sample Filters

path, and it will continue along the true branch when expanding the others.

Along the left-most path, the consumer immediately encounters an `end` node. This path is complete. Along the other paths, the consumer, following the true branch, immediately hits a `get`. The `get` matches the suspended `put` of the producer, so execution suspends at the consumer and resumes at the producer along both paths. Figure 5 gives the flow graph at this point.

The producer must now resume execution by exploring all possible control paths from its suspended `put`. Control continues to switch back and forth until no more progress can be made. A configuration labels each new node. Prior to adding a new node, its configuration is checked against the nodes already in the new graph—upon a match, the existing node is used rather than the new node.

Figure 6 shows the graph resulting from this composition. Filter Fusion is not finished at this point, however. Some paths reach a `get` without a corre-

sponding `put`. These paths are removed from the control flow, since they make no sense. Trimming often creates a conditional for which only one branch remains—in these cases, we may remove the conditional too. In general, trimming conditionals is an unsafe optimization. If, however, filters are properly composed such that `put`'s must always reach `get`'s (as they do here), the optimization can be both safe and effective. The nodes to be safely trimmed have double borders in Figure 6. Figure 7 gives the trimmed graph.

The final step of Filter Fusion is transforming the matched `put`'s and `get`'s into assignments to and reads from a temporary, respectively. The temporary is unique to a particular filter composition. Each suspended `put` that is copied into the composition graph becomes a write to the temporary, and all `get`'s become reads of the temporary. Figure 8 highlights the transformed nodes in the resulting graph with double borders.

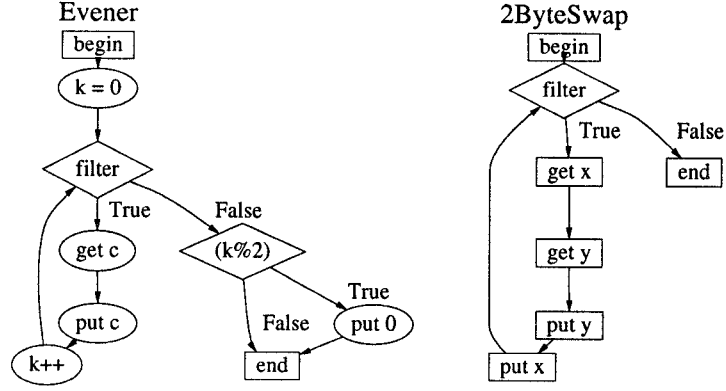


Figure 3: Original Control Flow Graphs




Figure 4: Stage 1

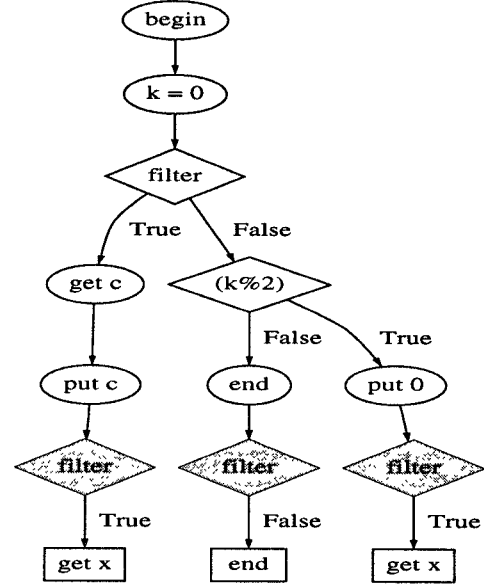


Figure 5: Stage 2

5 Algorithm

Filter Fusion is done pairwise, starting with the source and its immediate consumer. Because the composition of a source and a general filter is itself a source, this method can compose arbitrarily many filters. (Filter Fusion can operate in the opposite direction too, but giving the less general algorithm here is simpler.)

FFC implements Filter Fusion with a work-list algorithm. Elements of the work-list represent configurations that have been added to the control-flow graph (CFG), but whose successors have not. The algorithm is responsible for computing the successors and adding them to the CFG and the work-list, when necessary. No computed configuration already in the

CFG will be added to the work-list, since the previous instance can be reused in its place. This ensures termination. It also bounds number of nodes in the fused graph by the product of the number of nodes in the input graphs. (In practice, the code size will not increase to this maximum, particularly when merging filters with the same size data units.)

Figure 9 gives the algorithm. Let x be a CFG node. Its configuration is defined by $x.orig[producer]$, $x.orig[consumer]$, and $x.tag$. $x.orig[producer]$ and $x.orig[consumer]$ represent the last nodes visited in the two filters when this node was generated. $x.tag$ indicates which filter generated this node. Additional attributes of x , $insn$ and $successors$, denote the node's actual instruction and its CFG successors.

The algorithm begins by adding a *start* configu-

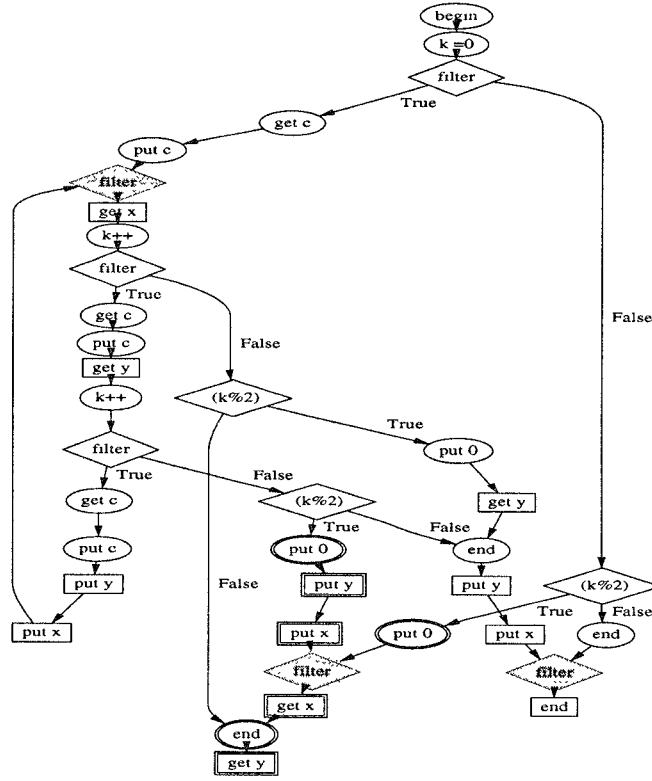


Figure 6: Untrimmed Control Graph

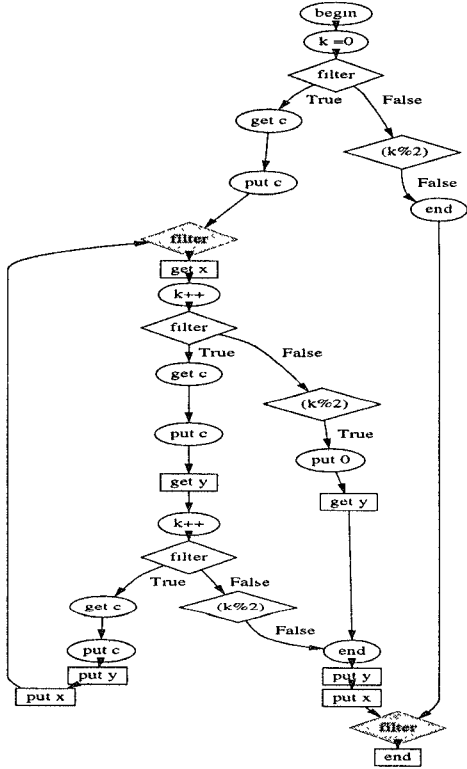


Figure 7: Final Control Graph




Figure 8: Assignment Substitution

```

Procedure Fusion()
  start.orig[producer] := producer's start node      // Initialize start node's configuration
  start.orig[consumer] := consumer's start node
  start.tag := producer
  start.insn := empty instruction
  CFG := { start }                                  // Seed CFG and worklist.
  worklist := { start }
  repeat
    x := Pop(worklist)
    if x.insn  $\notin$  trigger[x.tag] then              // {put,end} for producer; {get,end} for consumer.
      this := x.tag                                  // Stay with current filter.
      other := not x.tag
    else                                              // Switch to other filter.
      this := not x.tag
      other := x.tag
    endif
     $\forall i \in x.orig[this].successors$  do              // Follow all paths.
      node := new node
      node.orig[this] := i                          // Store current nodes.
      node.orig[other] := x.orig[other]
      node.tag := this                               // Tag which filter derived node.
      node.insn := i.insn
      if node  $\notin$  CFG then
        CFG := CFG  $\cup$  node
        Append(worklist, node)
        x.successors := x.successors  $\cup$  node
      else                                          // Reuse existing node.
        x.successors := x.successors  $\cup$  CFG[node]
      endif
    end  $\forall$ 
  until worklist =  $\phi$ 
end Fusion

```

Figure 9: Algorithm

ration that represents the initial nodes of each of the input graphs to both the CFG and the worklist. *start* will be the beginning node of the resulting graph. While elements remain in the work-list, they are removed one at a time, to compute their successors. Recall that successor nodes may or may not come from the same control flow graph as a node, *x*, itself (e.g., the successor of a **put** in the producer comes from the consumer, but the successor of a simple statement in the producer would also come from the producer). *trigger*[*producer*] represents the set of nodes that cause control to switch from the producer to consumer, **put** and **end**. *trigger*[*consumer*] is a set consisting only of **get**. “not tag” alternates between producer and consumer.

After computing the source of *x*'s successors, the algorithm simply follows the control flow from the

last executed statement in that source graph to find the actual successor instructions. Each successor has a configuration that is checked against the CFG to determine if it already exists. If the configuration already exists, the control flow arc out of *x* simply points to the existing configuration. If the configuration is new, it is added to the CFG and the work-list. The new configuration is also the target of the arc from *x*.

The algorithm describes the steps to compute the untrimmed graph. Trimming the graph of *dangling put* nodes is straightforward. Also, a little additional bookkeeping is necessary to transform **put**'s and **get**'s into assignment and reads of temporaries.

6 Experimental Results

FFC is a 200-line Icon program [GG90]. **FFC** is a preprocessor that generates C code from a compact specification language. We tested **FFC**'s code against modular and hand-integrated implementations on a variety of platforms and compilers. The modular implementation uses arrays for communicating values between adjacent filters. Because **gcc** consistently produced worse code than the vendor compilers, we aborted its use. (**gcc** had difficulty re-ordering basic blocks to avoid chains of jumps. It also did not handle copy propagation and dead-code elimination as well as the vendor compilers.)

To test **FFC**-generated code, we created different filter compositions. Our first test was the following composition.

ReadFromArray → **Evener** → **2ByteSwap** → **CRC32**
→ **WriteToArray**

These filters (1) read bytes from an array, (2) pad arrays to an even length, (3) swap bytes, (4) compute **CRC32** checksumming, and finally, (5) write the bytes to an array. Appendix A contains the specification for **CRC32** and the composition. Figure 10 gives the final flow graph. Note that some chunks of code are replicated multiple times and that the graph is quite complicated given the simple nature of its constituent filters. Table 1 shows the size of several fused filters on both a DEC Alpha and the Sun SPARCsystem 10. Since **FFC** may replicate the same code multiple times, the final fused filter may contain a great amount of C code. Although the C code produced by the Filter Fusion compiler was much larger than that of the modular and hand-integrated implementations, the object code sizes were very nearly comparable. Compiler optimizations eliminate much of the redundancy.

FFC-generated code must be optimized because of its heavy reliance on temporary variables and arbitrary control flow. The code particularly stresses—and finds deficiencies in—a compiler's copy propagation and dead code elimination optimizations. Unfortunately, in many cases, all of the available compilers failed to eliminate useless counters or to propagate copies. In addition, the compilers did not appear to unroll unstructured loops. Therefore, **FFC**-generated code's performance suffered. Performing these optimizations by hand yields code that performs comparably to that of the hand-integrated routines.

We timed three different implementations of the first five-filter composition: modular, hand-integrated, and **FFC**-generated integration. Table 2 gives the results of running these filters 10,000 times

over a 10,000 element array. All tests were run on four different architectures using the vendor's C compilers.

The second experiment again used five filters.

ReadFromArray → **RLE** → **2ByteSwap** → **PES** →
WriteToArray

These filters (1) read bytes from an array, (2) Run-Length Encode bytes, (3) swap bytes, (4) **PES** "encrypt" bytes, and (5) write the bytes to an array. Timings for 10,000 iterations over a 10,000 byte input yields the results in Table 3. Appendix A contains the specification for **PES** and **RLE**. The **PES** encryption simply takes two bytes and creates a new pair of bytes by alternating bits from the originals. This filter is a self inverse—the original data can be recovered from the encrypted data by simply reapplying the filter. The data-accessing overhead in most encryption methods (e.g., RSA, DES) is dwarfed by intensive computation, and therefore gains little from Filter Fusion. The **PES** filter is not intended to model an actual encryption method, but rather to model another lightweight data manipulation [Abb93].

Our third experiment used more new filters. We added a decryption phase, another byteswapping phase, and a decode phase to the previous experiment. To verify the correctness of **FFC**-generated code, we created the following composition.

ReadFromArray → **RLE** → **2ByteSwap** → **PES** → **PES**
→ **2ByteSwap** → **RLD** → **WriteToArray**

These filters (1) read bytes from an array, (2) Run-Length Encode bytes, (3) swap bytes, (4) **PES** encrypt bytes, (5) **PES** decrypt bytes, (6) swap bytes, (7) Run-Length Decode bytes, and finally, (8) write the bytes to an array. Timings for 10,000 iterations over a 10,000 byte input yields the results in Table 4. Appendix A contains the specification for **RLD** filter.

FFC-generated output is always superior to modular code. **FFC**-generated output typically is slower than hand-integrated code, but only because of the C compiler's shortcomings. Filter Fusion allows the programmer to maintain a modular design and implementation without sacrificing performance.

This exhaustive computation of all possible execution paths is tedious and error-prone when done by hand. Fortunately, **FFC**, an implementation of Filter Fusion, automates this transformation. Filter Fusion allows the programmer to forget about this complex work, and focus on optimizing independent filters in a modular fashion.

Program	C size (in lines)	Alpha Binary Size (in bytes)	Sparc Binary Size (in bytes)
Modular Implementation	36	2,784	1,955
Hand Integrated	28	2,592	1,898
Filter Fusion	197	2,976	2,323
Fused & Tuned	144	3,040	2,127

Table 1: Code Size: ReadFromArray → Evener → 2ByteSwap → CRC32 → WriteToArray

Architecture	Fusion Technique (in sec.)		
	No Integration	Hand Integration	Filter Fusion
DEC/Alpha	21.9	8.3	9.1
Sun/Sparc	26.2	12.0	13.9
HP/700	38.5	19.9	28.2
Mips R2000A	66.7	33.9	41.4

Table 2: ReadFromArray → Evener → 2ByteSwap → CRC32 → WriteToArray

Architecture	Fusion Technique (in sec.)		
	No Integration	Hand Integration	Filter Fusion
DEC/Alpha	22.3	11.9	12.2
Sun/Sparc	26.7	13.1	13.5
HP/700	37.0	26.2	30.2
Mips R2000A	64.6	37.6	41.8

Table 3: ReadFromArray → RLE → 2ByteSwap → PES → WriteToArray

Architecture	Fusion Technique (in sec.)		
	No Integration	Hand Integration	Filter Fusion
DEC/Alpha	41.7	20.1	21.4
Sun/Sparc	47.2	26.3	33.2
HP/700	68.4	41.3	52.3
Mips R2000A	121.5	70.1	88.4

Table 4: ReadFromArray → RLE → 2ByteSwap → PES → PES → 2ByteSwap → RLD → WriteToArray

References

- [Abb93] Mark B. Abbott. *A Language-Based Approach to Protocol Implementation*. PhD thesis, University of Arizona, 1993.
- [ASU86] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, Reading, Massachusetts, 1986.
- [BD77] R. M. Burstall and John Darlington. A transformation system for developing recursive programs. *Journal of the ACM*, 24(1):44–67, January 1977.
- [CJRS89] David D. Clark, Van Jacobson, John Romkey, and Howard Salwen. An analysis of tcp processing overhead. *IEEE Communications Magazine*, June 1989.
- [CT90] David D. Clark and David L. Tennenhouse. Architectural considerations for a new generation of protocols. In *Proceedings of the SIGCOMM '90 Symposium*, pages 200–208, September 1990.

- [DAPP93] Peter Druschel, Mark B. Abbott, Michael A. Pagels, and Larry L. Peterson. Network subsystem design: A case for an integrated data path. *IEEE Network Magazine*, July 1993.
- [GG90] Ralph E. Griswold and Madge T. Griswold. *The Icon Programming Language*. Prentice Hall, 1990.
- [JGS93] Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, New York, 1993.
- [MMO⁺95] A. B. Montz, D. Mosberger, S. W. O'Malley, L. L. Peterson, and T. A. Proebsting. Scout: A communications-oriented operating system. In *Proceedings of the 5th Workshop on Hot Topics in Operating Systems*, pages 58–61. IEEE Computer Society Press, May 1995.
- [OPM94] Sean O'Malley, Todd A. Proebsting, and A. Brady Montz. USC: A universal stub compiler. In *Proceedings of SIGCOMM 94 Conference on Communications Architectures, Protocols and Applications*, pages 295–306, August 1994.
- [Wad84] Philip Wadler. Listlessness is better than laziness: Lazy evaluation and garbage collection at compile-time. In *Proceedings of the ACM Symposium on Lisp and Functional Programming*, pages 45–52, August 1984.
- [Wad90] Philip Wadler. Deforestation: Transforming programs to eliminate trees. *Theoretical Computer Science*, 73:231–248, 1990.

A Five-Filter Specification

The FFC specifications of the `Evener`, `2ByteSwap`, `ReadFromArray`, and `WriteToArray`. are given in the paper. The specification below describes the remaining filters used in this paper. Figure 10 shows the first composition's final flow graph.

```
// Experiment 1
Compose prodeven ← ReadFromArray Evener
Compose prodevenBS ← prodeven Byteswap
Compose prodevenBSCRC ← prodevenBS CRC32
Compose fulltest ← prodevenBSCRC WriteToArray
```

```
// Experiment 2
Compose prodRLE ← producer RLE
Compose BSmid ← prodRLE Byteswap
Compose RLencodemid ← BSmid PES
Compose RLencode ← RLencodemid consumer
```

```
// Experiment 3
Compose prodRLE ← producer RLE
Compose BSmid ← prodRLE Byteswap
Compose pesmid ← BSmid PES
Compose unpesmid ← pesmid PES
Compose RLencodemid ← unpesmid Byteswap
Compose RLencode ← RLencodemid RLD
Compose massive ← RLencode consumer
```

```
Filter CRC32
Decls
    unsigned long crc = 0;
    unsigned char idx;
    unsigned char CRC32temp;
Code
    while filter
        get CRC32temp
        idx = (CRC32temp ^ crc);
        idx &= 0xff;
        crc >>= 8;
        crc ^= crctable[idx];
        put CRC32temp
    endwhile
    put crc & 0xff
    put (crc >> 8) & 0xff
    put (crc >> 16) & 0xff
    put (crc >> 24) & 0xff
End-Filter
```

```
Filter RLD
Decls
    unsigned char item, count, number;
Code
    while filter
        get item
        get number
        count = 0;
        while (count < number)
            put item
            count++;
        endwhile
    endwhile
End-Filter
```

```
Filter RLE
Decls
    unsigned char current, next, count;
Code
    if filter
        get current
    endif
    while filter
```

```

        count = 1;
        get next
        while (filter && (current == next))
            count++;
            get next
        endwhile
        put current
        put count
        current = next;
    endwhile
End-Filter

Filter PES
Decls
    unsigned char pes1, pes2;
    unsigned char out1, out2;
Code
    while filter
        get pes1
        get pes2
        out1 = (pes1 & 0xaa) + (pes2 & 0x55);
        out2 = (pes2 & 0xaa) + (pes1 & 0x55);
        put out1
        put out2
    endwhile
End-Filter

```

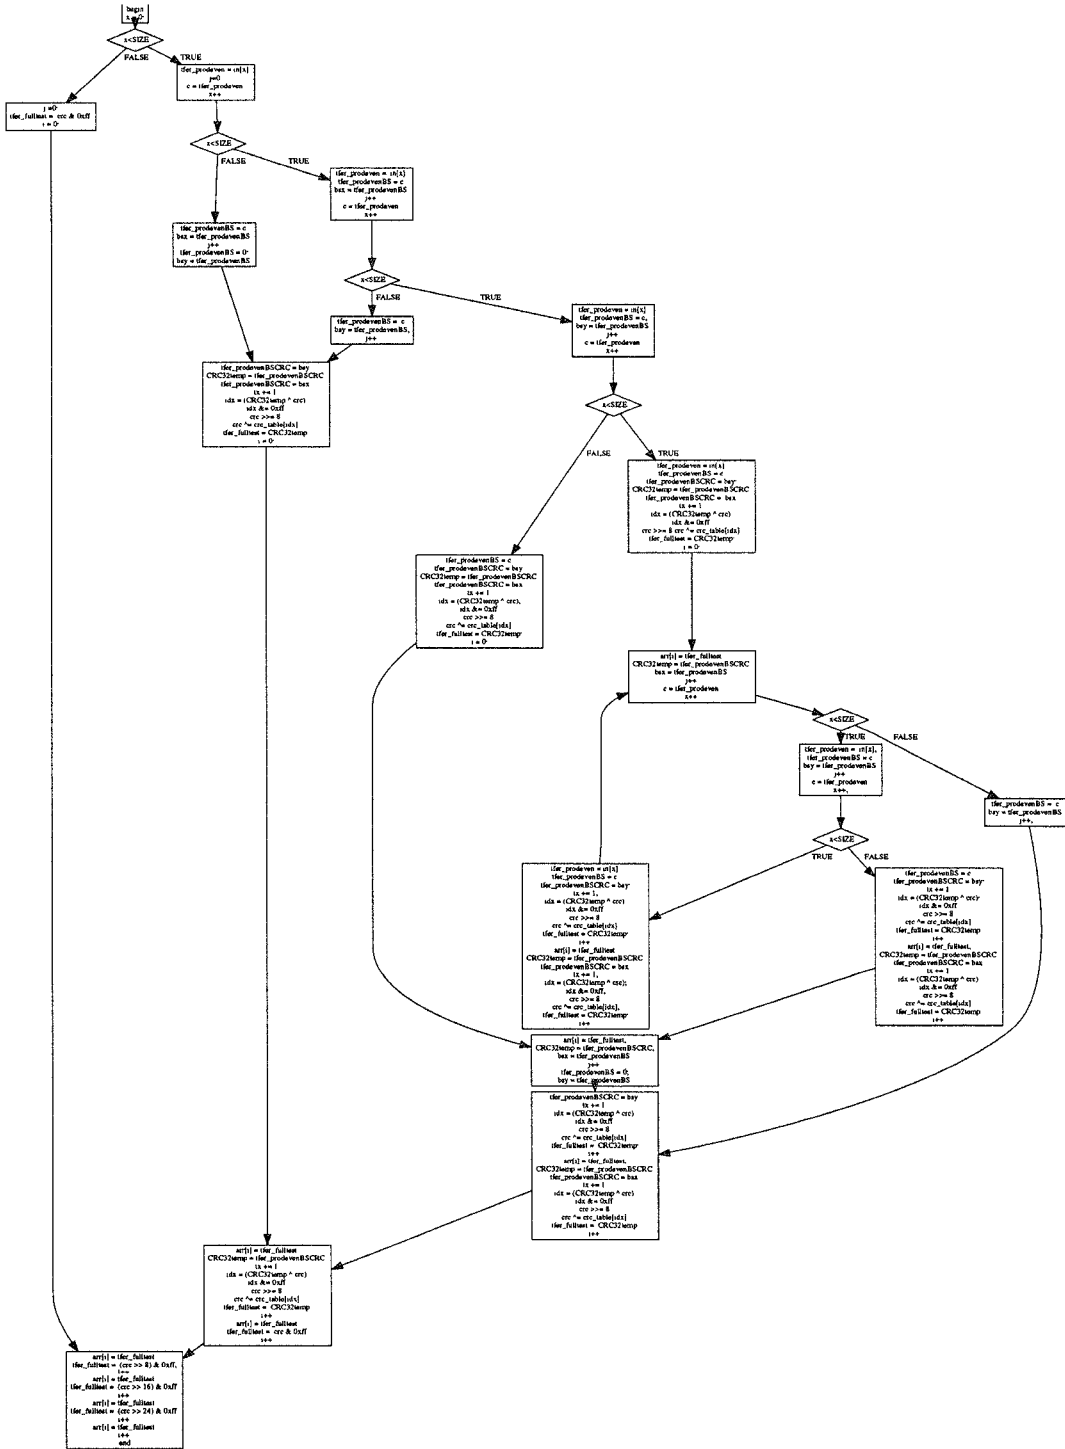


Figure 10: Final Composition: ReadFromArray → Evener → 2ByteSwap → CRC32 → WriteToArray. (Note the nonstructured control flow in the lower-right.)