# Enhancing *R* with Advanced Compilation Tools and Methods

## Duncan Temple Lang

*Abstract.* I describe an approach to compiling common idioms in *R* code directly to native machine code and illustrate it with several examples. Not only can this yield significant performance gains, but it allows us to use new approaches to computing in *R*. Importantly, the compilation requires no changes to *R* itself, but is done entirely via *R* packages. This allows others to experiment with different compilation strategies and even to define new domain-specific languages within *R*. We use the Low-Level Virtual Machine (*LLVM*) compiler toolkit to create the native code and perform sophisticated optimizations on the code. By adopting this widely used software within *R*, we leverage its ability to generate code for different platforms such as CPUs and GPUs, and will continue to benefit from its ongoing development. This approach potentially allows us to develop high-level *R* code that is also fast, that can be compiled to work with different data representations and sources, and that could even be run outside of *R*. The approach aims to both provide a compiler for a limited subset of the *R* language and also to enable *R* programmers to write other compilers. This is another approach to help us write high-level *descriptions* of what we want to compute, not *how*.

*Key words and phrases:* Programming language, efficient computation, compilation, extensible compiler toolkit.

## 1. BACKGROUND & MOTIVATION

Computing with data is in a very interesting period at present and this has significant implications for how we choose to go forward with our computing platforms and education in statistics and related fields. We are simultaneously (i) leveraging higher-level, interpreted languages such as *R*, MATLAB, Python and recently Julia, (ii) dealing with increasing volume and complexity of data, and (iii) exploiting, and coming to terms with, technologies for parallel computing including shared and nonshared multi-core processors and GPUs (Graphics Processing Units). These challenge us to innovate and significantly enhance our existing computing platforms and to develop new languages and systems so that we are able to meet not just tomorrow's needs, but those of the next decade.

*Duncan Temple Lang is Associate Professor, Department of Statistics, University of California at Davis, 4210 Math Sciences Building, Davis, California 95616, USA (e-mail: duncan@r-project.org).*

Statisticians play an important role in the "Big Data" surge, and therefore must pay attention to logistical and performance details of statistical computations that we could previously ignore. We need to think about how best to meet our own computing needs for the near future and also how to best be able to participate in multi-disciplinary efforts that require serious computing involving statistical ideas and methods. Are we best served with our own computing platform such as *R* (R Core Team, 2013)? Do we *need* our own system? Can we afford the luxury of our own system, given the limited resources our field has to develop, maintain and innovate with that system? Alternatively, would we be better off reimplementing *R* or an *R*-like environment on a system that is developed and supported by other larger communities, for example, *Python* or *Java*? Or can we leave it to others to build a new, fast computing environment that we can leverage within the field of statistics? Would "doing it right" give us an opportunity to escape some of the legacy in our current systems and position ourselves for more innovation? Would de-

veloping a new system splinter our already small community and reduce our effectiveness in disseminating new statistical methods as effectively as we do through *R*'s excellent package mechanism?

I have wrestled with these questions for over a decade. I don't believe there is a simple answer as to what is the best way to proceed. It is as much a social issue as a technical one. The *R* community is an amazing and valuable phenomenon. There is a large *R* code base of packages and scripts in widespread use for doing important science. Even if new systems do emerge and replace *R*, this will take several years. We need significant improvements in performance to make *R* competitive with other systems, at least for the near future. We must improve *R*'s performance to allow us to continue to deal with larger and more complex data and problems. In this paper, I discuss one direct approach to improve the performance of *R* code that is extensible and enables many people to further improve it.

The essence of the approach I am suggesting is conceptually quite simple and emerged in numerous other languages and platforms, around the same time we first started implementing it for *R*. The idea is that we compile *R* code directly to low-level native machine instructions that will run on a CPU or GPU or any device we can target. Instead of insisting that *R* code be evaluated by *the* one and only *R* interpreter, we may generate the code to perform the equivalent computations in a quite different way. We can dynamically compile fast native code by combining information about the code, the data and its representation being processed by that code, the available computing "hardware" (i.e., CPU or GPU or multi-core), the location of the different sources of the data, whether we need to handle missing values (NAs) or not, and so on. This is a form of just-in-time (JIT) compilation. It leverages additional knowledge about the context in which code will be run. It maps the code to low-level machine instructions rather than having it evaluated by a high-level interpreter.

The approach presented here is quite different from how programmers typically improve performance for *R* code. They manually implement the slow, computationally intensive parts in *C/C++*, and call these routines from *R*. I call this "programming around *R*." Instead, I am trying to "*compile* around *R*." In this approach, statisticians use familiar *R* idioms to express the desired computations, but the compiler infrastructure generates optimized instructions to realize the intended computations as efficiently as possible on the

hardware platform in use. The input from the statistician or analyst to this process is *R* code, not low-level *C/C++* code. This is good because humans can more easily understand, debug, adapt and extend code written in *R*. Furthermore, the compiler can "understand" what is intended in the high-level code and optimize in quite different ways. This also allows the code to be optimized in very different ways in the future. The high-level code says what to do, but not how to do it. How is left to the compiler.

What makes this approach feasible and practical now is the availability of the Low-Level Virtual Machine Compiler Infrastructure (*LLVM*) (Lattner and Adve, 2004). *LLVM* is the winner of the 2012 Association for Computing Machinery System Software Award (the same award conferred on the *S* language in 1999) and is a highly extensible compiler toolkit. LLVM is a *C++* library we can integrate into *R* (and other languages) to generate native code for various different CPUs, GPUs and other output targets. The ability to integrate this very adaptable and extensible tool into high-level languages is a "game changer" for developing compilation tools and strategies. We can use a technology that will continue to evolve and will be developed by domain experts. We can adapt these to our purposes with our domain knowledge. We do this within an extensible *R* package rather than in the *R* interpreter itself. This leaves the compilation infrastructure in "user" space, allowing development of any new compilation strategies to be shared without any changes to the *R* interpreter. This contrasts with *R*'s byte-code compiler and the byte-code interpreter which is part of the fixed *R* executable.

The ***Rllvm*** (Temple Lang, 2010b) package provides *R* bindings to the *LLVM C++* API. We can use this to generate arbitrary native code. The ***RLLVMCompile*** (Temple Lang and Buffalo, 2011) package provides a simple "compiler" that attempts to translate *R* code to native code by mapping the *R* code to instructions in *LLVM*, leaving *LLVM* to optimize the code and generate the native machine code.

Before we explore some examples of how we can use *LLVM* in *R* to improve performance and change our computational strategies for certain types of problems, it is worth thinking a little about the potential implications of fast *R* code:

- *Alternative data models.* On a practical level, if we can compile scalar (i.e., nonvectorized) *R* functions so that they are almost as fast as *C/C++* code, we can use them to process individual observations in

a streaming or updating manner. This means we can escape the highly-vectorized or "all data in memory" approaches that *R* strongly rewards.

- *Exporting code to alternative execution environments.* We can write *R* code and then export it to run in other different systems, for example, databases, *Python*, *JavaScript* and Web browsers. We can map the *R* code to *LLVM*'s intermediate representation (IR). We can then use emscripten (Zakai, 2010) to compile this directly to *JavaScript* code. For other systems, we can share the IR code from *R* and they can use their own *LLVM* bindings to compile it to native code for the particular hardware.

- *Richer data structures.* *R* provides a small number of primitive data types, for example, vectors, lists, functions. We can currently use these to create new aggregate or composite data structures. However, we can only introduce new and different data structures such as linked lists and suffix trees as opaque data types programmed in native (*C/C++*) code. When we compile code to native instructions, we also have the opportunity to have that new code use these different data structures and to represent the data differently from *R*. The same *R* code can be merged with descriptions of new data types to yield quite different native instructions that are better suited to particular problems.

- *Templating concepts.* Our ability to create native code from *R* code allows us to think about *R* functions or expressions differently. They are descriptions of what is to be done, without the specifics of how to do them. An *R* compiler can rewrite them or generate code that will behave differently from the *R* interpreter but give the same results (hopefully). The functions are "templates." The compiler can use knowledge of the particular representation of the data the functions will process to generate the native code in a more intelligent manner. For example, the code may access elements of a two-dimensional data set—rows and columns. There are two very different representations of this in *R*— data frames and matrices. How the individual elements are accessed for each is very different. The compiler can generate specialized code for each of these and might even change the order of the computations to improve efficiency (cache coherency) for these representations. The function is not tied to a particular data representation.

In summary, compiling *R* programs through LLVM yields novel computational potentials that are directly relevant to improving statistical learning and communication in the big data era. Compiling high-level code to native code is used in many systems. *Julia* is an interesting modern project doing this. NumPy (Jones et al., 2001) in *Python* is another. Several years ago, Ross Ihaka and I explored using *LISP* (Ihaka and Temple Lang, 2008) as the platform for a new statistical computing platform. The same ideas have been used there for many years and the performance gains are very impressive.

A very important premise underlying the approach in this paper is that the *R* project and its large code base are important, and will be for at least another 5 years. Users are not likely to immediately change to a new system, even if it is technically superior. For that reason, it is important to improve the performance of *R* now. It is also important to allow developers outside of the core *R* development team to contribute to this effort and to avoid many forked/parallel projects. For this, we need an extensible system within *R*, and not one that requires continual changes to the centralized *R* source code.

In addition to focusing on the immediate and near-term future and improving *R*, we also need to be exploring new language and computing paradigms within the field of statistics. *Julia* is an interesting modern project doing this. We need to foster more experiments so new ideas emerge. To do this, we also need to increase the quantity and quality of computing within our curricula.

Section 2 constitutes the majority of this paper. In it I explore different examples of computing in *R* and how compiling code can make the computations more efficient, both by simply obtaining faster execution speed and also by allowing us to change how we approach the problem. In Section 3 I discuss some additional general strategies we can exploit to improve computations in the future. I briefly discuss other exciting research projects to improve *R*'s efficiency and contrast them with the *LLVM* approach. I outline in Section 6 a road map of the ongoing work on the *LLVM* approach and other related projects as part of our future activities. The aim of this is to illustrate the feasibility of the entire approach.

In this paper I focus on reasonably standard *R* code and approaches that can be improved by generating native code. The semantics of the generated code are the same as the original *R* code. The approach also allows us to develop new languages and semantics within the compilation framework and to explore different computational models. However, the examples discussed in this paper stay within *R*'s existing computational model

in order to anchor the discussion and avoid too many degrees of freedom becoming a distraction. I do hope that we will explore new semantics and language features within *R* via compilation.

## 2. ENHANCING *R* WITH ADVANCED COMPILATION TOOLS AND METHODS

In this section we'll explore some examples of how we can write code in *R* and compile it to machine code. These explore different strategies and illustrate how we can approach computations differently when we have the option to compile code rather than only interpret it. We have chosen the problems for several reasons. They are each reasonably simple to state, and they illustrate the potential benefits of compilation. Like most benchmarks, some of the examples may not reflect typical use cases or how we would do things in *R*. However, most of these problems are very concrete and practical, and represent ways we would like to be able to program in *R*, were it not for the performance issues. In this way, the examples illustrate how we can continue to use *R* with an additional computational model and can overcome some of the interpreter's performance issues while still using essentially the same *R* code.

NOTE. In the following subsections, we present absolute and comparative timings for different approaches and implementations to the different tasks. These timings were performed on three different machines. We used a MacBook Pro running OS X (2.66 GHz Intel Core i7, 8 GB 1067 MHz DDR3) and also two Linux machines. The first of these Linux machines is an older 2.8 GHz AMD Opteron, 16 GB. The second is a much more recent and faster machine—3.50 GHz Intel Core i7-2700K, with 32 GB of RAM. Additionally, the different machines have different compilers used to compile *R* itself and these may impact the timings. We used GCC 4.2.1 on OS X, and GCC 4.3.4 on the first Linux machine and both GCC 4.8.0 and clang on the second Linux box. In all cases, *R* was compiled with the -O3 optimization level flag. The absolute times are quite different across machines, as we expect, and the within-machine relative performance of our *LLVM* generated code to native code differs between OS X and Linux. However, the within-machine results are very similar across Linux machines. Finally, our current steps to optimize the native code we generate with *LLVM* are quite simple and we expect to improve these in the near future.

We have not included the time to compile the code in our measurements. There are two steps in this—compiling the *R* code to intermediate representation (IR) and then compiling the IR to native code. The former can be done once, and the latter for each *R* session and is done in *LLVM*'s *C++* code. There are several reasons for omitting these steps in the timings. First, our focus is on tasks that take a long time to run in *R*, for example, many hours or days. Compilation time will be on the order of, at most, a few minutes and so the compilation time is negligible. Second, we expect that the compiled code will be reused in multiple calls and so the overhead of compiling will be amortized across calls. We have also ignored the time to byte-compile *R* functions, or compile and install *C/C++* code to be used in *R* packages.

### 2.1 The Fibonacci Sequence

The Fibonacci sequence is an interesting mathematical sequence of integers defined by the recurrence/recursive relation

$$F_n = F_{n-1} + F_{n-2}, \quad n \geq 0$$

with $F_0 = 0$ and $F_1 = 1$. We can implement this as an *R* function in an easy and obvious manner as

```
fib = function(n)
{
  if (n < 2L)
    n
  else
    fib(n - 1L) + fib(n - 2L)
}
```

For simplicity, we don't verify that `n` is nonnegative, assuming the caller will provide meaningful inputs. This maps the mathematical description of the sequence to a computational form essentially in a one-to-one manner. This is a good thing, as it makes the code and computations easy to understand, debug and maintain. However, this is a scalar function and not vectorized, meaning that it computes the value of the Fibonacci sequence for a single integer value rather than element-wise for a vector of inputs. This makes it slow in *R* if we want to compute multiple values from the sequence, for example, apply it to each element of a vector. Instead of implementing the function in this natural form, to gain performance, we would look to other implementations. Since the sequence is described by a recurrence relationship, there is a simple closed-form formula for computing the *n*th value of the sequence which can be easily implemented in *R* as a vectorized

function of *n*. Alternatively, we might use memoization to remember results computed in previous calls to the function to avoid repeating computations. We might even use a lookup table of pre-computed results for common input values, or some combination of these approaches. The key is that to get good performance, we have to think about the problem quite differently. Instead, we'll explore how we make the simple implementation above perform better and hope to avoid having to change our entire approach or rely on *R*'s other vectorized operations.

We use the function *compileFunction()* in the **RLLVMCompile** package to create a native compiled version of the *fib()* function with

```
fib.ll = compileFunction(fib,
  Int32Type, list(n = Int32Type))
```

We have to specify the type of the return value and also the type of the input(s), that is, n in this case. For this function, both the return type and the input are regular integer values corresponding to the 32-bit integer type Int32Type. We could use a 64-bit integer by using the type Int64Type if we wanted to deal with larger numbers. In fact, we can create two separate and different versions of this function with different types with two calls to *compileFunction()*. This is a simple illustration of how easy it is to adapt the same *R* code to different situations and create different compiled routines with different characteristics.

*compileFunction()* can return an *R* function which we can invoke directly. However, by default, it currently returns an object representing the compiled routine in *LLVM*. We can invoke the routine using this object and the *.llvm()* function, analogous to the *.Call()* and *.C()* functions in *R*. So

```
.llvm(fib.ll, 30)
```

calls our compiled routine and returns the value 832,040. Unlike the *.Call()/.C()* functions, the *.llvm()* function knows the expected types of the routine's parameters and so coerces the inputs to the types expected by the routine. In this case, it converts the *R* value 30 from what is a numeric value to an integer.

After verifying that the routine gives the correct results, we can explore the performance of the code. This recursive function is very computationally intensive. When calculating, for example, fib(30), we calculate fib(28) twice [once for each of fib(30 - 1) and fib(30 - 2)] and, similarly, we compute fib(27) multiple times and so on. This repetition is one of the reasons the code is so slow. We'll compare the time to evaluate fib(30) using three different versions of the *fib()* function: the original interpreted function, the *LLVM*-compiled routine (fib.ll) and a version of *fib()* that is compiled by *R*'s byte-compiler. The *LLVM*-compiled routine is the fastest. Table 1 shows the elapsed times for each and a ratio of the time for each function relative to the time for the interpreted function. These convey the relative speedup factor. We see that on a Macbook Pro, the *LLVM*-compiled routine is 600 times faster than the *R* interpreter. On a Linux machine, the speedup is a bit smaller, but still very significant at a factor of 500. While we have attempted to reduce the variability of these timings, we have observed different speedups ranging from 400 to 600 on OS X and 230 to 540 on Linux. The timings we report here are the most recent (rather than the "best").

Again, this is a simple example and not necessarily very representative of how we would calculate the Fibonacci sequence in production code. However, the ability to express an algorithm in its natural mathematical form makes it easier to program, verify and extend. We would very much like to be able to write code in this manner, without sacrificing good run-time performance.

TABLE 1
*Timings for computing Fibonacci Sequence Values*

|  | OS X | | Linux 1 | | Linux 2 | |
| --- | --- | --- | --- | --- | --- | --- |
|  | **Time** | **Speedup** | **Time** | **Speedup** | **Time** | **Speedup** |
| Interpreted *R* code | 80.49 | 1.00 | 112.70 | 1.0 | 51.780 | 1.0 |
| Byte-compiled *R* code | 31.70 | 2.53 | 45.85 | 2.5 | 21.620 | 2.4 |
| ***Rllvm*-compiled code** | 0.12 | 653.90 | 0.21 | 526.4 | 0.097 | 531.0 |

These are the timings for a call to fib(30) using the regular *R* function, the byte-compiled version and the *LLVM*-compiled version. To improve the accuracy of the timings, we calculate the duration for 20 replications for the two slower functions and 200 replications of the *LLVM*-compiled routine and divided the duration by 10. The *LLVM*-compiled version is clearly much faster.

## 2.2 2-Dimensional Random Walk

Ross Ihaka developed a very instructive example of writing straightforward *R* code compared with clever, highly vectorized *R* code as a means to illustrate profiling in *R* and how to make code efficient. The task is simulating a two-dimensional random walk. It is very natural to write this as a loop with N iterations corresponding to the N steps of the walk. For each step, we toss a coin to determine whether we move horizontally or vertically. Given that choice, we toss another coin to determine whether to move left or right, or up or down. Then we calculate and store the new location. We'll call this the naïve approach and the code is shown in Figure 1. After several refinements based on profiling and nontrivial knowledge of *R*, Ihaka defines a very efficient *R* implementation of the random walk, shown in Figure 2. It removes the explicit loop, samples all N steps in one call to *sample()*, and determines the positions using two calls to the *cumsum()* function. This makes very good use of several of *R*'s vectorized functions which are implemented in *C* code and therefore fast.

We manually compiled the naïve implementation using ***Rllvm*** and, similarly, used *R*'s byte-compiler to create two compiled versions of this function. We then simulated a 10 million step random walk using each of the original naïve function, the byte-code compiled function, the fully vectorized version and the *LLVM*-compiled version. Table 2 shows the relative speedups.

```
rw2d1 =
function(n = 100)
{
  xpos = ypos = numeric(n)
  truefalse = c(TRUE, FALSE)
  plusminus1 = c(1, -1)
  for(i in 2:n) {
       # Decide whether we are moving
       #  horizontally or vertically.
    if (sample(truefalse, 1)) {
       xpos[i] = xpos[i-1] +
                    sample(plusminus1, 1)
       ypos[i] = ypos[i-1]
     }
     else {
       xpos[i] = xpos[i-1]
       ypos[i] = ypos[i-1] +
                    sample(plusminus1, 1)
     }
   }
   list(x = xpos, y = ypos)
}
```

FIG. 1. *The naïve implementation of the* 2-*D random walk.*

```
rw2d5 =
# Sample from 4 directions, separately.
function(n = 100000)
{
  xsteps = c(-1,  1,  0,  0)
  ysteps = c( 0,  0, -1,  1)
  dir = sample(1:4, n - 1, replace = TRUE)
  xpos = c(0, cumsum(xsteps[dir]))
  ypos = c(0, cumsum(ysteps[dir]))
  list(x = xpos, y = ypos)
}
```

FIG. 2. *The fast*, *vectorized implementation of the* 2-*D random walk.*

We see that the manually vectorized *R* function is 175 times faster than the naïve implementation, illustrating how important vectorization is to make *R* code efficient. However, we also see that compiling the naïve implementation with *LLVM* outperforms even the vectorized version, taking about between 55% to 65% of the time of the vectorized version. This is probably due to the compiled code using a single loop, while the vectorized version has two calls to *cumsum()* and hence at least one additional *C*-level loop over the N steps.

### 2.3 Sampling a Text File

Suppose we have one or more large comma-separated value (CSV) files. For example, we can download airline traffic delay data for each year as an approximately 650 megabyte CSV file from the Research and Innovative Technology Administration (RITA), part of the Bureau of Transportation, at http://www.transtats.bts.gov/DL_SelectFields.asp?Table_ID=236. Rather than working with the entire data set, we might choose to take a random sample of the observations. (We don't concern ourselves here with the appropriateness of a simple random sample.) We'll also assume that we know the number of observations in the CSV file.

How do we efficiently extract a sample of the lines from the file? We could use UNIX shell tools, but it is difficult to randomly generate and specify the lines to sample. Sampling the indices is something we want to do in *R*, but then passing all of these to a shell command is awkward, at best. Alternatively, we could do the entire sampling in *R*. We could read the entire file into memory [via the *readLines()* function] and then subset the ones we want. However, this requires a significant amount of memory. We first store all of the lines, then make a copy of the ones we want and then discard the larger vector. This may not be feasible, as we may not have enough memory, or it may simply be too slow.

TABLE 2
*Timings for simulating a 2-D Random Walk*

| | OS X | | Linux 1 | | Linux 2 | |
|---|---|---|---|---|---|---|
| | Time | Speedup | Time | Speedup | Time | Speedup |
| Interpeted *R* code | 171.08 | 1.0 | 196.6 | 1.0 | 100.3 | 1.0 |
| Byte compiled code | 123.92 | 1.4 | 120.8 | 1.6 | 60.51 | 1.66 |
| Vectorized *R* code | 0.97 | 176.5 | 1.8 | 106.8 | 0.63 | 159.46 |
| ***Rllvm***-compiled code | 0.52 | 329.3 | 1.1 | 180.3 | 0.40 | 250.12 |

We generate 10 million steps for each approach. We compare a manually vectorized implementation in *R* code with a naïve version written in *R*, both a byte-compiled and *LLVM*-compiled version of that naïve function. The vectorized version is 175 times faster than the regular *R* function. However, the *LLVM*-compiled version outperforms the vectorized version, most likely by removing one *C*-level loop.

We can think of different strategies. One is to first identify the indices of all of the lines we want in our sample, and then read the file in chunks until we get to a line that is in our sample. We store that line and continue to read from where we are up to the next line in our sample, and so on. To make this work, we need to be able to continue to read from where we currently are in the file. We can use an *R* file connection to do this.

Our first step is to generate the vector of the line numbers we want to sample using, for example,

```
lineNum = sort(sample(1:N,
  sampleSize))
```

where `N` is the number of lines in the CSV file. We have sorted the line numbers, as we will read the sample lines in the file sequentially. The next step is to determine how many lines there are between successive lines in our sample. We can compute this in *R* with

```
lineSkips = diff(c(0, lineNum))
```

which gives a vector of the pairwise difference between successive elements. For example, suppose the first two lines we want to sample are 60 and 200. The first two elements in `lineSkips` will be 60 and 140. We can then read the first two lines in our sample with

```
con = file("2012.csv", "r")
readLines(con, 60)[60]
readLines(con, 140)[140]
```

Each element of `lineSkips` tells us how many lines to read to get the next line in our sample. So next we need a function that can read that many lines and return the last of these. The following function does this:

```
readTo = function(numLines, con)
  readLines(con, numLines)[numLines]
```

The final step to obtaining our entire sample is to call *readTo()* for each element of `lineSkips`, for example,

```
readSelectedLines =
  function(lineSkip, file)
  sapply(lineSkip, readTo, file)
```

To obtain our sample, we call *readSelectedLines()*, passing it the variable `lineSkips` and our open connection:

```
con = file("2012.csv", "r")
sample =
  readSelectedLines(lineSkips, con)
```

Each of these functions is concise and efficient since *sapply()* is essentially implemented as a *C*-level loop within the *R* interpreter. Using the connection and *readLines()* to read blocks of lines in *readTo()* is efficient as it uses *C* code within *R*. Unfortunately, it does involve reading, allocating, storing, subsetting and discarding a potentially large character vector returned by each call to *readLines()*. However, we only want a single line at the end of that vector in each call. While each call involves significantly fewer lines than reading the entire file, allocating a large character vector still slows the computations, as it extensively involves the memory manager in *R*. A different approach to avoid the memory issue is to change the *readTo()* function so that it reads each line individually and then returns the last one. We could change it to

```
readTo =
function(numLines, con)
{
  ans = ""
  for(i in 1:numLines)
    ans = readLines(con, 1)
```

```
    ans
}
```

Again, this is straightforward and easy to understand. Unfortunately, it is extremely slow as we are now looping in *R* over almost every line in the file.

The idea of reading one line at a time would work well if we could avoid the overhead of the *R* loop mechanism. We can do this if we compile this new version of *readTo()* into native code. We can almost do this now, but we need to have an equivalent of *readLines()* to read a single line of a file. This is exactly what the standard *C* routine *fgets()* does. Similar to a connection, we pass *fgets()* a pointer to an opaque *C*-level *FILE* data structure, and it puts the contents of the next line it reads into a location in memory that we also provide. For simplicity of exposition, we will define our own function *Fgets()* in *R* as a proxy to call *fgets()* with

```
Fgets = function(file)
  fgets(ptr, 1000L, file)
```

This is *R* code and it just assumes there is a function named *fgets()* and that `ptr` is somehow (the address of) an array in memory with 1000 character elements, that is, space for a long string. We won't run this code in *R*, so these variables [`ptr`, `file` and *fgets()*] don't actually have to exist in *R*. Instead, we will allocate them in *LLVM* for the compiled, native routine we generate from *Fgets()*.

We compile the *Fgets()* function in an *LLVM* module, a collection of routines and variables, using *compileFunction()*. We also define the module-level "global" variable *ptr* to be the pointer to the array we want, after creating the actual array of 1000 characters as another global variable. When compiling *Fgets()*, we also need to tell the compiler about the signature of the external *fgets()* routine so that it can make the call to *fgets()* correctly. We do this via

```
mod = Module()
FILEType = pointerType(Int32Type)
declareFunction(list(StringType,
  StringType, Int32Type, FILEType),
  "fgets", mod)
```

[While we have done this explicitly, we could automate this step using the **RCIndex** (Temple Lang, 2010a) package to obtain the signature programmatically.] We also need to tell the *LLVM* run-time engine how to locate the *fgets()* routine which we do with

```
llvmAddSymbol("fgets")
```

Note that in our *Fgets()* function, we assumed that the longest line was less than 1000 characters. We can specify a different length if we knew or suspected otherwise. Similarly, we didn't provide any error checking about whether we had reached the end of the file. This is because we are assuming that the caller knows the total number of lines and is sampling only up to, at most, that number. This is an example of the context-specific shortcuts we can make when compiling the code for a particular situation and not writing general, robust code which can be used in many different situations. We could also tell the compiler to add these tests for us, if we wanted, but can avoid the extra computations when we know they are redundant.

How do we obtain the instance of the *FILE* data type to pass to the compiled *Fgets()* routine? We can use the *C* routine *fopen()* and again, we can write an *R* function that mimics that and then compile it. However, the **RLLVMCompiler** package has a function to automate the creation of that proxy function in *R*, if we know the signature of the *C* routine of interest. So this example illustrates how we can dynamically create bindings to existing compiled routines in different libraries. In the case of *FILE*, we can also use the existing function *CFILE()* in the **RCurl** (Temple Lang, 2002) package.

So now we can read a single line from an open *FILE* object in *R* via our compiled *Fgets()* routine. We can redefine our *readTo()* function as

```
readTo =
function(numLines, con)
{
  ans = ""
  for(i in 1:numLines)
    ans = Fgets(con)
  ans
}
```

This is almost identical to the original function above but replaces the call `readLines(con, 1)` with `Fgets(con)`. Now we can compile this into native code via *compileFunction()* and the resulting code will be quite fast.

We now have a fast replacement for reading up to the next line in our sample. The last step is to make *readSelectedLines()* fast. Recall that this was implemented simply as `sapply(lineSkip, readTo, file)`. When we compile this as returning an *R* character vector, our compiler recognizes the *sapply()* call and converts this into a loop in native code and populates and returns a new *R* character vector.

TABLE 3
*Timings for sampling a CSV file*

| | OS X | | Linux 1 | | Linux 2 | |
|---|---|---|---|---|---|---|
| | Time | Speedup | Time | Speedup | Time | Speedup |
| Interpreted *R* loop & *readLines()* | 68.93 | 1.0 | 103.25 | 1.0 | 42.78 | 1.0 |
| ***Rllvm***-compiled loop & *Fgets()* | 3.278 | 21.09 | 6.54 | 15.8 | 2.59 | 16.5 |
| *C* code (***FastCSVSample***) | 3.0 | 22.97 | 6.28 | 16.4 | 2.40 | 17.8 |

We use vectorized code in *R* to read blocks of data and extract the final line of each block. The *LLVM* approach compiles simple *R* functions that read one line at a time. The ***FastCSVSample*** does the same thing with manually written *C* code. The compiled approaches avoid the memory usage related to *readLines()* and see a nontrivial speedup. The *C*-code in the ***FastCSVSample*** package outperforms the *LLVM*-compiled version, but both approaches outperform the approach using *R*'s connections and *readLines()* functionality which are also implemented with *C* code.

In summary, we have compiled three *R* functions [*Fgets()*, *readTo()* and *readSelectedLines()*] and these now allow us to read one line at a time and use the minimal amount of memory to collect the lines for our sample, but using two loops in native code rather than in *R*.

We can now compare the performance of our *R*-based approach using *readLines()* to consume chunks of lines and our compiled version that reads one line at a time. In addition to these two approaches, we also have a manual *C* implementation essentially equivalent to our *LLVM*-compiled approach in the ***FastCSVSample*** package (Temple Lang, 2013). Our timings are based on extracting a sample of one hundred thousand lines uniformly from a CSV file that contains one hundred million lines—the same lines for each approach. The elapsed times are given in Table 3. We see that our compiled approach of reading one line at a time is around twenty times faster than collecting many unnecessary lines with *readLines()* and looping in *R*, even with *sapply()*. The difference between the *LLVM* and native *C* approaches may be inherent, but also possibly due to different optimization techniques that we may be able to enhance with *LLVM*. In short, we can outperform *R*'s native vectorized code by compiling our relatively straightforward *R* code.

The exposition of this example may make it seem more complicated than it is. Essentially, we want to efficiently read one line of a file at a time in order to get to the next line in our sample. We compiled the *Fgets()* function for this and then compiled two other functions in *R* to perform loops over the number of lines. The important implications from this example is that we can sidestep *R*'s memory management, get fine-grained control over computations using dynamically generated routines, and we can use existing native routines and data structures, such as *fgets()* and *FILE*, in our *R* code that will be compiled. We could already dynamically call native routines directly from *R* using, for example, **rdyncall** (Adler, 2012) or **Rffi** (Temple Lang, 2011). What is important here is that we are also compiling the iterations and not doing them in *R*.

### 2.4 Fusing Loops

Consider computing the log-likelihood for a given vector of observations x and a density function, say, *dnorm()*. In *R*, we can write the log-likelihood very efficiently as

```
sum(log(dnorm(x, mu, sigma)))
```

Indeed, we could reduce this to sum(dnorm(x, mu, sigma, log = TRUE)), but the purpose of this example is to consider a general sequence of calls to vectorized *R* functions.

Each of the functions *dnorm()*, *log()* and *sum()* are built-in to *R* and are implemented in *C* code, and two of them use the very efficient .*Primitive()* mechanism. As a result, this code seems to be as fast as it can be. This is true given the way *R* interprets the expression, one sub-expression at a time. However, there are two ways we can make this more efficient by compiling such an expression. The first is by reducing the number of loops (in *C*) from three to one. Generally, if we have $n$ nested-calls to element-wise and vectorized functions, we can reduce $n$ loops to one. Second, we can typically eliminate at least one allocation of a potentially large vector. A third way we might speed up the computations is to use parallel capabilities such as a GPU or multiple cores or CPUs. We won't discuss this here, but it is conceptually quite straightforward to do generally when we are compiling the code dynamically. Indeed, the ability to programmatically combine a particular function with a general parallel strategy makes it more expedient than writing it ourselves in *C*/*C++*.

How does *R* evaluate the expression above? It uses three separate loops. Ignoring pedantic details, essentially *R* evaluates the call to *dnorm()* and so loops over all of the elements of x and computes the density at each of those values. It stores these values in a newly allocated vector and then returns that. This becomes the input to the call to *log()*. *R* then iterates over the elements of this vector and computes the *log()* for each individual value. In this case, *R* may recognize that it doesn't need to create a new vector in which to return the results, but that it can reuse the input vector since it is essentially anonymous. The final step in the overall expression is the call to *sum()* and this iterates over the elements of the vector it receives and returns a single scalar value.

Importantly, there are three loops over three vectors all of the same length, and we allocate one new and large vector. We could use a single loop and avoid allocating this intermediate vector by rewriting the code as

```
normalLogLik =
function(x, mu = 0, sigma = 1) {
  ans = 0
  for(val in x)
    ans = ans +
      log(dnorm(val, mu, sigma))
  ans
}
```

Instead of the vectorized calls in *R*, we have put scalar function calls inside a single loop. We have combined the calls to *dnorm()* and *log()* together. Then we took the result for each element and added it to the cumulative sum. This combination of operations is called loop fusion and for large vectors can yield significant performance improvements.

This new scalar version is faster by avoiding the loop and allocation. Of course, it is evaluated in *R* and so will be much slower. We could write this in *C*, but it would be very specific to the log-likelihood for a Normal density. Generally, we would have to write implementations for various sequences of calls, for example, for different density functions [i.e., sum(log(pdf(x, ...)))], and expressions involving other functions [e.g., prod(dchisq(x^2, p))]. This isn't practical. However, given our ability to dynamically generate native code, we can compile any expression such as our original expression sum(log(dnorm(x, mu, sigma))) into the native equivalent of our scalar code above.

To compile the *normalLogLik()* function above, we need to be able to call scalar versions of the *log()* and *dnorm()* routines. The *log()* function is available in the ubiquitous math library (**libm**) and we can just refer to it. The Normal density function is not standard. We can arrange for our native code to invoke *R*'s *dnorm()* function for each scalar value in the vector. This is both awkward and inefficient. Instead, we can write our own version of *dnorm()* directly in *R*. While this would be slow to invoke many times in *R*, we will compile our *dnorm()* and *normalLogLik()* functions together into a single module and both will be fast. Another possible approach, in this case, is to take advantage of the good design and modularity of the Rmath library. It provides the routine *dnorm4()* as a regular native routine (unconnected with *R*'s data types, etc.) and so we can invoke it, just as we do the *log()* routine.

For reasons that are not quite clear at present, on the OS X machine, our loop-fused version takes about 40% longer than the *R* code for 10 million observations and 27% longer for 100 million observations. Again, we suspect that we will be able to improve the *LLVM*-compiled code by exploring more of its optimization facilities. However, on the Linux machines, we do see a speedup, even for 10 million observations where the *LLVM* loop-fused code runs in about 75% the time of the *R* code. The timings and relative performances are given in Table 4. Regardless of the exact numbers, the results indicate that compiling our own code is competitive with manually writing vectorized routines in *R*, and that we can outperform these built-in *C* routines.

A difference between the two approaches is that *R* uses the *.Primitive()* mechanism rather than a standard function call which we have to do via the *.llvm()* function. However, not only do we reduce three loops to one, but we also avoid dealing with missing values (*NA*s) and additional parameters such as base for the *log()* function. So we should be doing even better. If we have access to multiple cores or GPUs, we may be able to execute this code much more efficiently simply via parallel execution. By fusing the loops operations together, we can also avoid three separate transitions from the host to the GPU and transferring memory between the two systems more times than we need.

We explicitly wrote the *normalLogLik()* function to show how to fuse the loops. We could also have written the original expression sum(log(dnorm(x, mu, sigma))) as

```
Reduce('+', Map(log, Map(dnorm, x,
  MoreArgs = list(mu, sigma))))
```

TABLE 4
*Times and relative performance for fusing loops*

|  | OS X | Linux 1 | Linux 2 |
|---|---|---|---|
| ***Rllvm***-compiled fused loops | 0.73 | 1.69 | 0.84 |
| Interpreted *R* vectorized functions | 0.52 | 2.28 | 1.07 |
| Regular *R* time/***Rllvm*** time | 0.71 | 1.35 | 1.27 |

The first two rows show the times for fusing the loops by compiling the *R* with *LLVM* and using a sequence of calls to *R*'s vectorized functions. The final row shows the ratio of the two times within each machine. Fusing the loops is slower on OS X, but faster on Linux.

By explicitly using these functional programming concepts, it is easy for us to see how to fuse loops and rewrite the code into the loop above. The **RLLVMCompile** package can recognize such an expression and compile it to the loop-fused instructions. We can either require *R* programmers to do this in order to gain the performance from native code or we can try to make the compiler recognize the vectorized nested function call idiom of the form `f(g(h(x)))`.

## 2.5 Computing Distances between Observations

Distances between pairs of observations are important in common statistical techniques such as clustering, multi-dimensional scaling, support vector machines and many methods that use the "kernel trick" (Schölkopf and Smola, 2001). *R* provides the *dist()* function that allows us to compute the distance between all pairs of observations in a matrix or data frame, using any of six different metrics. The core computations are implemented in *C* and are fast. However, there are some issues and rigidities.

The *dist()* function insists that the data passed to the *C* code are represented as a matrix, and so will make a copy of the data if a data frame is given by the caller. For large data sets, this can be a significant issue as we will essentially have two copies of the data in memory. Also, the *dist()* function only accepts a single data set and computes the distances between all pairs of observations within it. In contrast, a reasonably common situation is that we start with two separate data sets—X and Y—and want to compute the distance between each observation in X and each observation in Y, but not the distances between pairs of observations within X or within Y. Not only do we risk having three copies of the data in memory (the two separate data frames, the two combined into one data frame and then converted to a matrix), but the *dist()* function will also perform many unnecessary computations for these within-same-set obser-

vations that we will discard. If we have two data sets with $n1$ and $n2$ observations, respectively, the *dist()* function computes $(n1 + n2) \times (n1 + n2 - 1)/2$ distances. We are only interested in $n1 \times n2$ of these. As $n1$ and $n2$ diverge, the number of unnecessary computations increases, and this is especially burdensome if the number of variables for each observation is large.

Another rigidity is that the choice of distance metric is fixed. If we wanted to introduce a new distance metric, it would be useful to be able to reuse the *C* code underlying *dist()*. We could do this with a function pointer in *C*, but the code for *dist()* would need to be modified to support this. Accordingly, if we want to introduce a new metric, we have to copy or re-implement the entire *C* code.

The *C* code underlying *dist()* can use parallel capabilities (OpenMP) if they are detected when *R* is compiled. We cannot use GPUs or change the parallel strategy within an *R* session without rewriting the *C* code. As a result, we would like to be able to express the computations in *R* and select a different strategy for parallelizing the computations at run-time.

In short, as useful as *dist()* is, we would like it to be much more flexible. We want to be able to compute the distances between two sets of observations, not within a single data set; use a data frame or a matrix or perhaps some other data representation without making a copy of the data; introduce new metrics within the same infrastructure; and use different parallel computing approaches. The current *dist()* function in *R* cannot help us meet these goals and is essentially static/fixed code.

The package **pdist** (Wong, 2013) provides a way to compute pairwise distances between two data sets. This avoids the redundant computations. Unfortunately, it only supports the Euclidean metric and also insists on matrices being passed to the *C* code. Also, it has no support for parallel computing.

If we could write the basics of the *dist()* function in *R* and make it fast, we could address all of the en-

hancements we listed above as well as make the code more comprehensible and accessible to users. The basic approach to computing the distance between each pair of observations in two data sets X and Y can be expressed in *R* with the following quite specific/rigid function (written to aid compiling):

```
dist =
function(X, Y, nx = nrow(X),
  ny = nrow(Y), p = ncol(X))
{
  ans = numeric(nx * ny)
  ctr = 1L
  for(i in 1:nx) {
    for(j in 1:ny) {
      total = 0.0
      posX = i
      posY = j
      for(k in 1:p) {
        total = total +
          (X[posX] - Y[posY])^2
        posX = posX + nx
        posY = posY + ny
      }
      ans[ctr] = sqrt(total)
      ctr = ctr + 1L
    }
  }
  ans
}
```

The basic steps are to loop over each observation in the first data set (X) and then to loop over each observation in the other data set (Y). For each pair of observations, we compute the distance between them via the third nested loop. We could have made this simpler (and more general) by using a vectorized *R* expression or calling a function to do this final loop. However, we have *inlined* the computations directly for a reason. Suppose we had written this part of the computation as (X[i,] - Y[j,])^2. Unfortunately, in *R*, this would cause us to create two new intermediate vectors, one for each of the specific rows in the two data sets. This is because the row of each data set is not a simple vector containing the elements of interest which we can pass to the subtraction function (-). Instead, we have to arrange the data in each row of the matrix or data frame into a new vector of contiguous values. This is where *R* is convenient, but inefficient. This does not happen in the *C* code for *R*'s builtin *dist()* routine, or ours, as it uses matrices and knows how to access the elements

individually rather than creating a new temporary vector. We use this same approach in our loop. We also could allocate the vectors for the row values just once and reuse them for each observation, but we still have to populate them for each different observation.

To avoid the intermediate vectors, our code explicitly accesses the individual elements X[i, k] and Y[j, k] directly. A matrix in *R* is merely a vector with the elements of the matrix arranged sequentially in column order. Therefore, the first element of observation *i* in X is at position *i* in the vector. The second element of the *i*th observation is at position i + nrow(X), and so on. To compute the distance between the two observations, we loop over the *p* variables present in each of the observations and compute the difference.

The code illustrates these computations for the Euclidean distance. We could easily change this to implement other distance metrics. We could do this by changing the code either manually or programmatically by replacing the expression (X[posX]-Y[posY])^2 with, for example, abs(X[posX] - Y[posY]). Rewriting code programmatically is a powerful feature that allows us treat *R* code as a template.

We can compile this three-level nested loop *R* code via **RLLVMCompile** to native instructions. Our compiler currently works primarily with primitive data types and has limited support for working directly with *R* objects, for example, knowing the dimensions of an *R* matrix. Accordingly, we arrange to pass the matrices and their dimensions to the routine and currently have to explicitly specify the signature:

```
distc = compileFunction(dist,
  REALSXPType,
  list(X = DoublePtrType,
    Y = DoublePtrType,
    nx = Int32Type, ny = Int32Type,
    p = Int32Type))
```

In the future, we will allow the caller to specify just the two data sets (X and Y). However, we are making the representation as matrices more explicit here, which is valuable information for the compiler.

Now that we have the native code, we can then compare this to using *R* code that computes the same distances but does so by combining the two data sets, calls *dist()*, and converts the result to the sub-matrix of interest. This comparison favors our code since this is the form of our inputs and the expected form of the output. However, these are quite reasonable. We timed the functions to compute the distances for two data

TABLE 5
*Timings for computing pair-wise distances*

|                                           | OS X  | Linux 1 | Linux 2 |
| ----------------------------------------- | ----- | ------- | ------- |
| **Rllvm**-compiled code                   | 8.72  | 11.94   | 6.22    |
| *R dist( )* function (calling native code) | 14.74 | 79.65   | 27.37   |
| Speedup factor                            | 1.69  | 6.67    | 4.4     |

This shows the total elapsed time for distance computations with 40 variables and 8000 and 1000 observations in the two data sets. In the *R* approach with the *dist( )* function, there is extra memory allocation and also 80% of the distances computed are discarded. We outperform the native *R* implementation on both platforms.

sets of size 8000 and 1000 observations, each with 40 variables. In this case, 80% of the distances computed using *R*'s *dist( )* function are irrelevant and discarded. Table 5 shows the results and illustrates that by doing fewer computations, we do indeed outperform the native *C* code in *R*, on both platforms. If we had used data sets with similar numbers of observations, the results would have been less dramatic. However, with 3000 observations in Y, the *LLVM*-generated native code was still three times faster on Linux and only 18% slower on OS X.

Comparing the results above to similar native code in the **pdist** package, the timings again show that native *C* code in **pdist** outperforms our *LLVM*-compiled code, 60% faster on one machine and 9 times faster on another. This illustrates that there is room for significant improvement in our *LLVM* compilation. However, the fact that we can outperform *R*'s native approach is encouraging. That we can readily adapt this to different purposes and different computational strategies indicates significant opportunities and potential.

As a final note, we could remove the third loop and insert a call to a function to compute the distance for these two variables, for example, `euclidean(X[i,], Y[j,])`. The compiler could recognize that X and Y are matrices and arrange for the compiled version of the *euclidean()* function to access the elements as we have displayed above, that is, without computing the intermediate vector for each row. If we tell the compiler X and/or Y are data frames, it would generate different code to access the elements so as to avoid these intermediate vectors. Since the compiler has the opportunity to compile both the code for the main loop and for the metric function together and knows the representations of the inputs, it can create better code than if we wrote these separately and more rigidly.

## 3. POSSIBLE COMPILATION ENHANCEMENT STRATEGIES

The examples in the previous section explored different ways we could change the way we compute in *R* with new facilities for generating native code. We considered compiling *R* code to native routines, reusing existing native routines within these generated routines, and changing the computational strategies we employ within *R* to embrace these new approaches. There are many other simple examples we could consider to improve the performance of *R* code. One is the ability to write functions that focus on scalar operations and then to create vectorized versions of these automatically. Given a scalar function *f()*, we can write a vectorized version as `sapply(x, f, ...)` or with *mapply()*. The compiler can then turn this into a native loop. Indeed, many of the performance gains are achieved by making looping faster. They also potentially reduce the necessity to use vectorized code in *R* and so hopefully make programming in *R* more intuitive for new users.

In addition to handling loops, there are several other aspects of *R*'s evaluation model that we might be able to improve by choosing different compilation strategies in different contexts. The idea is that the *R* user compiling the code may have more information about the computations, the data and its representation, or the available computing resources than the compiler does by examining the code. This extra information is important. The programmer may be able to give hints to a compiler, or choose a different compiler function/implementation altogether, to control how the code is understood and the native instructions are generated. The following are some reasonably obvious and general improvements we might be able to infer or make in certain situations. Guided by the *R* user, different compilers may yield different code, and even different semantics, for the same code.

### 3.1 Omitting Checks for *NA* Values

Many of the *C* routines in *R* loop over the elements of a vector and must check each element to see if it is a missing value (*NA*). This code is general purpose code and so this test is a fixed part of almost every computation involving that routine. However, when we dynamically generate the code, we may know that there are no missing values in the data set on which we will run that code and so omit the code to perform these additional, redundant tests.

Similarly, in our example of sampling a CSV file, we knew the number of lines in a file and we knew that each call to the *fgets()* routine would succeed. As a result, we did not have to check the return status of the call for reading at the end of the file. We also assumed that the largest line was less than 1000 characters and didn't validate this in each iteration. The same applies when we are accessing elements of a vector as to whether we first need to check that the index is within the extent of the array or not, that is, bounds checking. When we can verify this conceptually (within a loop over a vector), or by declaration by the user, we can omit these checks.

These tests are typically simple and not computationally expensive. However, they can become significant when the instructions are invoked very often, for example, in a loop over elements of a large vector.

### 3.2 Memory Allocation

In our example discussing loop fusion, we saw that not only could we reduce the number of overall iterations in a computation, but we also could reduce memory usage. We avoided creating a vector for the result of the call to *dnorm()* [and *log()*]. There are potential opportunities to further reduce memory usage.

*R* uses the concept of pass-by-value in calls to functions. In theory, *R* makes a copy of each argument in a call to a function. (Lazy evaluation means that some arguments are never evaluated and so not copied.) However, the *R* interpreter is smarter than this and only copies the object when it is modified, and only if it is not part of another object. When compiling *R* code, we want to be able to determine that an object is not modified and avoid copying it. By analyzing code, we can detect whether parameters can be considered read-only and so reduce memory consumption in cases where *R* cannot verify that it is safe to avoid copying an object. We can identify this within regular *R* code, however, we would have to modify the interpreter to make use of this information. When generating native code with,

for example, *compileFunction()*, we can make use of this information dynamically, bypassing the *R* interpreter.

Another example where we can reduce the memory footprint of code is when we can reuse the same memory from a different computation. For example, consider a simple bootstrap computation something like the following *R* pseudo-code:

```
for(i in 1:B) {
  d.star = data[sample(1:n, n,
    replace = TRUE), ]
  ans[[i]] = T(d.star, ...)
}
```

In *R*'s computational model, we will allocate a new data frame d.star for each bootstrap sample. This is unnecessary. We can reuse the same memory for each sample, as each sample has the same structure and only differs in the values in each cell. By analyzing the sequence of commands rather than executing each one separately without knowledge of the others, we can take advantage of this opportunity to reuse the memory. We can also reuse the same vector to store the result of the repeated calls to *sample()*. It is reasonably clear that this is what we would do if we wrote this code in *C*, reusing the same data structure instances. However, this is not possible within *R*, as the individual computations are not as connected as they are in the large-picture *C* code. When we dynamically generate native code, we can utilize this large-scale information.

Similarly, some *R* scripts create a large object, perform several computations on it and then move to other tasks. Code analysis can allow us to identify that the object is no longer being used and so we can insert calls to remove the object. However, we may be able to recognize that the object is no longer needed, but that subsequent tasks can reuse the same data format and representation. In that case, we can reuse the memory or at least parts of it.

### 3.3 Data Representations

The small number of fundamental data types in *R* makes computational reasoning quite simple, both to use and to implement. Of course, the choice of data type and structure can be important for many computations. Sequences, for example, 1:n or seq(along = x), are common in *R* code and these are represented in *R* as explicit vectors containing all of the values in the sequence. We have seen that we can avoid creating the sequence vector and populating it when it is used as a

loop counter. Similarly, we can represent a regular sequence with the start, end and stride, that is, the increment between elements. When generating the code, we then access elements of such a sequence using appropriate calculations specialized to that sequence type.

In many cases, *R*'s simple data types cause us to use an integer when we only need a byte, or even just a few bits, to represent a few possible values/states. The **snpStats** (Clayton, 2011) package does this successfully using bytes to reduce the memory footprint for large genomic data. Again, the operations to subset data in this different format need to be modified from the default. Doing this element-wise in *R* is excessively slow. However, when we generate native code, we are free to use different ways to access the individual elements. This idea is important. We specify what to do in the code, but not precisely how to do it. When generating the code, we combine the code and information about how to represent the data and generate different code strategies and realizations. This is somewhat similar to template functions in *C++*, but more dynamic due to run-time compilation/generation with more contextual information.

There are several other aspects of *R* code that we can compile, for example, matching named arguments at compile time rather than at run time.

## 4. OVERVIEW OF GENERATING CODE WITH *LLVM*

In this section we will briefly describe the basic ideas of how we generate code with *LLVM*, **Rllvm** and **RLLVMCompile**. This is a little more technical and low level than our examples and readers do not need to understand this material to understand the main ideas of this paper or to use the compiler or the compiled code. We are describing it here to illustrate how other *R* programmers can readily experiment with these tools to generate code in different ways.

We'll use the Fibonacci sequence and the *fib()* function example again, as it illustrates a few different aspects of generating code.

Our *fib()* function in *R* expects an integer value and returns an integer. The body consists of a single *if–else* expression. This contains a condition to test and two blocks of code, one of which will be evaluated depending on the outcome of that condition. To map this code to *LLVM* concepts, we need to create different instruction blocks, each of which contains one or more instructions. When we call the routine, the evaluation starts in the first instruction block and executes each of its instructions sequentially. The end of each instruction block has a terminator which identifies the next

block to which to jump, or returns from the routine. Jumping between blocks allows us to implement conditional branching, loops, etc.

For our *fib()* function, we start with an entry block that might create any local variables for the computations. In our function, this block simply contains code to evaluate the condition n < 2 and, depending on the value of this test, the instruction to branch to one of two other blocks corresponding to the expressions in the *if* and the *else* parts. In the *if* block (i.e., *n* is less than 2), we add a single instruction to return the value of the variable n. In the block corresponding to the *else* part, we add several low-level instructions. We start by computing n - 1L and then call *fib()* with that value and store the result in a local variable. Then we calculate n - 2L, call *fib()* and store that result. Then we add these two local intermediate results and store the result. Finally we return that result. Figure 3 shows the code in what is called Intermediate Representation (IR) form that *LLVM* uses. This illustrates the low-level computations.

While the code for this function is reasonably simple, there are many details involved in generating the native code, such as defining the routine and its parameters, creating the instruction blocks, loading and storing values, and creating instructions to perform subtraction, call the *fib()* function and return a value. The *LLVM C++* API (Application Programming Interface)

```
; ModuleID = 'fib'

define i32 @fib(i32 %n) {
entry:
  %0 = icmp slt i32 %n, 2
  br i1 %0, label %"body.n < 2L", label %body.last

"body.n < 2L":                      ; preds = %entry
  ret i32 %n

body.last:                          ; preds = %entry
  %"n - 1L" = sub i32 %n, 1
  %1 = call i32 @fib(i32 %"n - 1L")
  %"n - 2L" = sub i32 %n, 2
  %2 = call i32 @fib(i32 %"n - 2L")
  %"fib(n - 1L) + fib(n - 2L)" = add i32 %1, %2
  ret i32 %"fib(n - 1L) + fib(n - 2L)"
}
```

FIG. 3. *Intermediate Representation for the compiled fib() routine. We create the Function and Block objects and create and insert LLVM instruction objects corresponding to the expressions and sub-expressions in the R function. The result is this low-level description in intermediate form which LLVM can optimize and compile to native code for different targets, for example, a CPU or GPU. The different blocks have a label (e.g.,* body.last) *and correspond to different parts of an if statement or possibly parts of a loop, generally.*

provides numerous classes and methods that allow us to create instances of these conceptual items such as *Function*s, *Block*s, many different types of instructions and so on. The **Rllvm** package provides an *R* interface to these *C++* classes and methods and allows us to create and manipulate these objects directly within *R*. For example, the following code shows how we can define the function, the entry instruction block and generate the call `fib(n - 1)`:

```
mod = Module()
f = Function("fib", Int32Type,
  list(n = Int32Type), mod)
start = Block(f)
ir = IRBuilder(start)
parms = getParameters(f)
n.minus.1 = binOp(ir, Sub, parms$n,
  createConstant(ir, 1L))
createCall(ir, f, n.minus.1)
```

We don't want to write this code manually ourselves in *R*, although **Rllvm** enables us to do so. Instead, we want to programmatically transform the *R* code in the *fib()* function to create the *LLVM* objects. The **RLLVMCompile** package does this. Since *R* functions are regular *R* objects which we can query and manipulate directly in *R*, we can traverse the expressions in the body of a function, analyze each one and perform a simple-minded translation from *R* concepts to *LLVM* concepts. This is the basic way the *compileFunction()* generates the code, using customizable handler functions for the different types of expressions. These recognize calls to functions, accessing global variables, arithmetic operations, if statements, loops and so on. They use the functions in **Rllvm** to create the corresponding *LLVM* objects and instructions.

Once the compiler has finished defining the instructions for our routine, *LLVM* has a description of what we want to do in the form of these blocks and instructions. This description is in this intermediate representation (IR). We can look at this "code" and it will look something similar to that shown in Figure 3. The IR code shows the somewhat low-level details of the blocks and instructions as we described above. We see the three blocks labeled `entry`, `body.n < 2L` and `body.last`. Again, it is not important to understand these details to be able to use the compiled routine translated from the *R* function. I show it here to illustrate the different steps in the compilation process and to indicate that an *R* programmer can chose to change any of these steps.

Next, we instruct *LLVM* to verify and optimize the code. At this point, we can call the new routine via the *.llvm()* function in **Rllvm** which corresponds to a method in the *LLVM* API. The first time the code is used, *LLVM* generates the native code from the IR form.

## 5. CONTRASTS WITH RELATED RESEARCH

There have been several projects exploring how to improve the performance of *R* code. We discuss some of these in this section.

*Byte-Compiler*: One of the most visible projects is the byte-code compiler developed by Luke Tierney (Tierney, 2001). This consists of an *R* package that provides the compiler and some support in the core *R* interpreter to execute the resulting byte-compiled code. The compiler maps the *R* code to instructions in the same spirit as *LLVM*'s instructions and intermediate representation. These instructions are at a higher-level than *LLVM*'s and are more specific to *R*.

The typical speedup provided by the byte-compiler is a factor of between 2 and 5, with much larger speedups on some problems. This may not be sufficient to obviate the need for writing code in *C/C++*. We probably need to see a factor of more than 10 and closer to 100 for common tasks.

The byte-code compiler is written in *R* and so others can adapt and extend it. However, the details of how the resulting byte-code it generates is evaluated is tightly embedded in the *C*-language implementation of the *R* engine. This means that if one wants to change the byte-code interpreter, one has to modify the *R* interpreter itself. While one can do this with a private version of *R*, one cannot make these changes available to others without them also compiling a modified version of *R*. In other words, the byte-code interpreter is not extensible at run-time or by regular *R* users. Furthermore, the *R* core development team does not always greet suggested enhancements and patches with enthusiasm. Therefore, this approach tends to be the work of one person and so has limited resources.

*Ra JIT Compiler*: The Ra extension to *R* and the associated **jit** package is another approach to using JIT (Just-in-Time) compilation in *R*. This focuses on compiling loops and arithmetic expressions in loops. Like the byte-compiler above, Ra requires almost no change to existing *R* code by *R* users—only the call to the function *jit()* before evaluating the code. The performance gain on some problems can be apparently as high as a factor of 27. (See http://www.milbo.

users.sonic.net/ra/times9.html.) Unfortunately, this is no longer maintained on CRAN, the primary central repository of *R* packages. This approach suffers from the fact that it requires a modified version of the *R* interpreter, again compiled from the *C*-level source. This places a burden on the author of Ra to continually update Ra as *R* itself changes. Also, it requires users trust Ra and take the time to build the relevant binary installations of Ra.

As we mentioned previously, important motivating goals in our work are to avoid modifying *R* itself, to allow other people to build on and adapt our tools, and to directly leverage the ongoing work of domain experts in compiler technology by integrating their tools to perform the compilation. Our approach differs from both the byte-compiler and Ra in these respects.

*R on the Java Virtual Machine*: There are several projects working on developing an implementation of *R* using the *Java* programming language and virtual machine. One is FastR (https://github.com/allr/fastr), which is being developed by a collaboration between researchers at Purdue, Oracle and INRIA. Another is Renjin (https://code.google.com/p/renjin/). Having *R* run on the *Java* virtual machine offers several benefits. There are many interesting large-data projects implemented in *Java*, for example, Apache's Hadoop and Mahout. Integrating *R* code and such projects and their functionality would be much tighter and effective if they are all on the same platform and share the same computational engine. Importantly, *R* would benefit from passively acquiring features in *Java* and its libraries, for example, security, threads. One very interesting development is that researchers in Oracle, collaborating with the developers of FastR, are creating a tool Graal (http://openjdk.java.net/projects/graal/) for compiling the code ordinarily interpreted by a high-level virtual machine, for example, FastR. This could yield the performance gains we seek, but by passively leveraging the general work of others merely by using widespread technologies. This contrasts with the ongoing development of *R* by a relatively small community and having to actively and manually import new technologies, features and ideas from other languages, systems and communities.

*Translating R code to C*: Another approach is to translate *R* code to *C*/*C*++ code. This is attractive, as it would give us similar speedup as we can get with *LLVM*, potentially produces human-readable code, and allows us to leverage the standard tools for these languages such as compilers, linkers and, importantly, debuggers. We can also potentially reuse (some of) the generated code outside of *R*. Simon Urbanek's **r2c** package (Urbanek, 2007) is an example of exploring translation of *R* code to *C* code.

The **Rcpp** (Eddelbuettel and François, 2011) (and **inline**) package are widely used in *R* to improve performance. The packages provide a way to include high-level *C*++ code within *R* code and to compile and call it within the *R* code. The *C*++ code uses an *R*-like syntax to make it relatively easier to write the *C*++ code. This has been a valuable addition to *R* to obtain performant code. The approaches that compile *R* code directly are preferable if they can get the same performance. The first reason is because the programmer does not have to program in *C*++. It is also harder for other programmers to read the code and understand what it does. The second reason is because the *C*++ code is essentially opaque to any *R* code analysis or compiler. If we do manage to generally compile *R* code effectively or implement an automated parallel computing strategy for *R* code, the *C*++ code cannot easily be part of this. For example, if we can map the *R* code to run as a GPU kernel on many cores, we cannot easily combine the *R* and *C*++ code to take advantage of these cores.

*Parallel Speedup*: There are several interesting projects that have aimed at improving the performance of *R* exclusively by running the code in parallel. This is very important and in some sense orthogonal to compilation of *R* code. If we speed up the computations on a single CPU, that speedup will benefit running code on each CPU. However, we also want to compile *R* code to take advantage of multiple CPU/GPUs. We hope to be able to integrate ideas from these projects into our compilation strategies. Unfortunately, some of them are no longer active projects, for example, pR and **taskPR**. This illustrates one aspect we have observed in the *R* community. Some researchers implement some ideas in *R*, sometimes as a PhD thesis, and then move on to other projects. One of the terrific aspects of *R* is the ongoing commitment to support the *R* community. This is probably a very significant reason for *R*'s widespread use and an important consideration when developing new environments and languages. It is one of the forces motivating our continued work within *R*, even if developing a new system would be intellectually more stimulating.

A very important aspect of all this work is to recognize that there are many positive ways to make *R* faster and more efficient. While one of these approaches *may* dominate others in the future, it is very important that

we should pursue comparative approaches and continue to motivate each other's work. There is much to be learned from these different approaches that will improve the others.

## 6. FUTURE WORK

Compiling (subsets of) *R* code and other Domain Specific Languages (DSLs) within *R* using *LLVM* is a promising approach that is certainly worth vigorously pursuing in the near term. The work is currently in its infancy—we started it in the summer of 2010, but have only recently returned to it after an almost three-year hiatus due to other projects (ours and other people's). However, the foundations of many of the important components are in place, that is, the ***Rllvm*** package, and the basics of the extensible and adaptable compiler mechanism in ***RLLVMCompile*** should allow us and others to make relatively quick progress, programming almost entirely in *R* to develop compilation strategies. However, there are many other tasks to do to make these transparent and reliable, and many related projects that will make them more powerful and convenient.

One of the immediate tasks we will undertake is to program some rich examples explicitly in *R* code. We are implementing *R* code versions of recursive partitioning trees, random forests and boosting. We also plan to explore compiling code for the Expectation Maximization (EM) algorithm and particle filters to run on GPUs. The aim is to share these sample *R* projects with the other researchers investigating different compilation strategies for *R* so that we compare approaches on substantive and real tasks we want to program in pure *R* code.

We plan to add some of the functionality available in *LLVM* that does not yet have bindings in ***Rllvm***. This includes topics such as different optimization passes and adding meta data to the instructions. We have also developed the initial infrastructure to compile *R* code as kernel routines that can be used on GPUs, that is, PTX (Parallel Thread Execution) code. Being able to generate kernel functions from *R* code, along with the existing *R*-CUDA bindings to manage memory and launch kernels from the host device, allows us to program GPUs directly within high-level *R* code. This contrasts with the low-level C code developed for existing *R* packages that target GPUs, for example, **gputools** (Buckner et al., 2009) and **rgpu** (Kempenaar and Dijkstra, 2010).

We will also be exploring different approaches to compiling the *R* code to run in parallel and distributed settings. We think that being able to use information about the distribution of the data to generate/compile the code will be important so that we can minimize the movement of data and keep the CPUs/GPUs busy on the actual computations rather than transferring the inputs and outputs to and from the computations.

Being able to write *R* code that directly calls *C* routines is very powerful. As we saw in relation to the *fgets()* routine in Section 2.3, we need to specify the signature for the routines we want to call. It is preferable to be able to programmatically identify these signatures rather than require *R* programmers to explicitly specify them. The ***RCIndex*** package is an *R* interface to **libclang** (Carruth et al., 2007), the parsing facilities for the clang compiler. This already allows us to read *C* and *C++* code in *R* and to identify the different elements it contains. This allows to not only determine the signatures of routines, but also discover different data structures, enumerated constants, etc. We can also go further and understand more about how the routines manipulate their arguments and whether they perform the memory management or leave it to the caller.

As we saw in each of our examples, information about the types of each parameter and local variable is a necessity to being able to compile using *LLVM*. Currently, the *R* programmer must specify this information not only for the function she is compiling, but also for all of the functions it calls. Again, we want to make this transparent, or at least only require the *R* programmer to specify this information when there is ambiguity. To this end, we are working on a type inference package for *R*. This starts with a known set of fundamental functions and their signatures. From this, we can determine the signatures of many higher level calls. As always, we cannot deal with many features of the language such as nonstandard evaluation, but we most likely can get much of the type information we need programmatically. Since *R*'s types are so flexible with different return types based on not only the types of the inputs, but also the content of the inputs, we need a flexible way to specify types. Perhaps the existing ***TypeInfo*** package (Temple Lang and Gentleman, 2005) or **lambda.r** package will help here. To analyze code for type information and for variable dependencies, we will build upon the ***CodeDepends*** (Temple Lang, Peng and Nolan, 2007) and **codetools** (Tierney, 2011) packages.

While these are some of the related activities we envisage working on, we also encourage others to collaborate with us or work independently using *LLVM* and optionally ***Rllvm*** and ***RLLVMCompile*** so that our community ends up with better tools.

## 7. CONCLUSION

We have described one approach to making some parts of the *R* language fast. We leverage the compiler toolkit infrastructure *LLVM* to generate native code. This allows us to incorporate technical knowledge from another community, both now and in the future. We can generate code for CPUs, GPUs and other targets. We can dynamically specialize *R* functions to different computational approaches, data representations and sources, and contextual knowledge, giving us a new and very flexible approach to thinking about high-level computing.

We are developing a simple but extensible and customizable compiler in *R* that can translate *R* code to native code. Not only does this make the code run fast, but it also allows us to compute in quite different ways than when we interpret the *R* code in the usual way. We can even outperform some of *R*'s own native code.

In no way should this work be considered a general compiler for all of the *R* language. There are many aspects of the *R* language we have not yet dealt with or considered. Vectorized subsetting, recycling, lazy evaluation and nonstandard evaluation are examples. We, or others, can add facilities to the compiler to support these when they make sense and are feasible.

The initial results from this simple approach are very encouraging. An important implication of this and other efforts to make *R* code efficient is that we can be benefit from writing high-level code that describes what to compute, not how. We then use smart interpreters or compilers to generate efficient code, simultaneously freeing *R* programmers to concentrate on their tasks and leveraging domain expertise for executing the code. We hope others will be able to use these basic building blocks to improve matters and also to explore quite different approaches and new languages within the *R* environment.

## ACKNOWLEDGMENTS

## SUPPLEMENTARY MATERIAL

The code for the examples in this paper, along with the timing results and their meta-data, are available from https://github.com/duncantl/RllvmTimings as a git repository. The versions of the ***Rllvm*** and ***RLLVMCompile*** packages involved in the timings can also be retrieved from their respective git repositories. The specific code used is associated with the git tag *StatSciPaper*.

## REFERENCES

ADLER, D. (2012). rdyncall: Improved foreign function interface (FFI) and dynamic bindings to C libraries. *R* package version 0.7.5.

BUCKNER, J., WILSON, J., SELIGMAN, M., ATHEY, B., WATSON, S. and MENG, F. (2009). The gputools package enables GPU computing in *R*. *Bioinformatics* **26** 135–135. Available at http://cran.r-project.org/web/packages/gputools/index.html.

CARRUTH, C., CHRISTOPHER, E., GREGOR, D., KOROBEYNIKOV, A., KREMENEK, T., MCCALL, J., ROSIER, C. and SMITH, R. (2007). libclang: *C/C++* translation unit parser library. Available at http://clang.llvm.org.

CLAYTON, D. (2011). snpstats: SnpMatrix and XSnpMatrix classes and methods. *R* package version 1.5.0.

EDDELBUETTEL, D. and FRANÇOIS, R. (2011). Rcpp: Seamless *R* and C++ integration. *Journal of Statistical Software* **40** 1–18.

IHAKA, R. and TEMPLE LANG, D. (2008). Back to the future: Lisp as a base for a statistical computing system. In *Proceedings in Computational Statistics* 21–33. Springer, Heidelberg.

JONES, E., OLIPHANT, T. and PETERSON, P. et al. (2001). SciPy: Open source scientific tools for Python.

KEMPENAAR, M. and DIJKSTRA, M. (2010). R/GPU: Using the Graphics Processing Unit to speedup bioinformatics analysis with *R*. Available at https://trac.nbic.nl/rgpu/.

LATTNER, C. and ADVE, V. (2004). LLVM: A compilation framework for lifelong program analysis and transformation. In *Proc. of the* 2004 *International Symposium on Code Generation and Optimization* (*CGO '04*), *San Jose*, *CA*, *USA* 75–88. IEEE Computer Society, Washington, DC. Available at http://llvm.org/.

R CORE TEAM (2013). *R: A Language and Environment for Statistical Computing*. *R* Foundation for Statistical Computing, Vienna, Austria.

SCHÖLKOPF, B. and SMOLA, A. (2001). *Learning with Kernels*. MIT Press, Cambridge, MA.

TEMPLE LANG, D. (2002). RCurl: General network (http/ftp/...) client interface for *R*. *R* package version 1.95-4.

TEMPLE LANG, D. (2010a). RCIndex: *R* interface to the clang parser's C API. Available at https://github.com/omegahat/RClangSimple.

TEMPLE LANG, D. (2010b). Rllvm: *R* interface to the Low-Level Virtual Machine API. Available at https://github.com/duncantl/Rllvm.

TEMPLE LANG, D. (2011). Rffi: Interface to libffi to dynamically invoke arbitrary compiled routines at run-time without compiled bindings. *R* package version 0.3-0.

TEMPLE LANG, D. (2013). FastCSVSample: An *R* package to sample lines from a text file. Available at https://github.com/duncantl/FastCSVSample.

TEMPLE LANG, D. and BUFFALO, V. (2011). RLLVMCompile: A simple LLVM-based compiler for *R* code. Available at https://github.com/duncantl/RLLVMCompile.

TEMPLE LANG, D. and GENTLEMAN, R. (2005). TypeInfo: Optional type specification prototype. *R* package version 1.27.0.

TEMPLE LANG, D., PENG, R. and NOLAN, D. (2007). CodeDepends: Analysis of *R* code for reproducible research and code comprehension. Available at https://github.com/duncantl/CodeDepends.

TIERNEY, L. (2001). Compiling *R*: A preliminary report. In *Proceedings of the* 2*nd International Workshop on Distributed Statistical Computing*, *March* 15–17, 2001 (K. Hornik and F. Leisch, eds.). Technische Universität Wien, Vienna, Austria.

TIERNEY, L. (2011). codetools: Code analysis tools for *R*. *R* package version 0.2-8.

URBANEK, S. (2007). *R* to C compiler. Available at http://www.rforge.net/r2c/index.html.

WONG, J. (2013). pdist: Partitioned distance function. *R* package version 1.2.

ZAKAI, A. (2010). Emscripten: An LLVM-to-JavaScript compiler. Available at https://github.com/kripken/emscripten/wiki.