

Introduction to High-Performance Computing with R

Dirk Eddelbuettel, Ph.D.

`Dirk.Eddelbuettel@R-Project.org`
`edd@debian.org`

Bank of Canada / Banque du Canada
December 22, 2008

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Motivation: Presentation Roadmap

We will start by *measuring* how we are doing before looking at ways to improve our computing performance.

We will look at *vectorisation*, as well as various ways to *compile code*. We will look at *debugging* tools and tricks as well.

We will discuss several ways to get more things done at the same time by using simple *parallel computing* approaches.

Next, we look at ways to compute with **R** *beyond the memory limits* imposed by the **R** engine.

Last but not least we look at ways to *automate* running **R** code.

Table of Contents

Motivation

Preliminaries

Measuring and profiling

Faster: Vectorisation and Compiled Code

Parallel execution: Explicitly and Implicitly

Out-of-memory processing

Automation and scripting

Summary

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Software Support

The tutorial is supported by a 'live cdrom'. The (updated) iso file `Quantian_Dec2008_tutorial.iso` can be downloaded from <http://quantian.alioth.debian.org/>.

The iso image contains a complete Debian operating system including the graphical KDE user interface. All the software demonstrated during the tutorial is available and fully functional. This includes

- ▶ R and all packages used,
- ▶ the accelerated Ra variant,
- ▶ Open MPI, NWS, Slurm and more,
- ▶ Emacs, ESS and a few other tools.

The versions correspond to the to the December 2008 snapshot of the upcoming Debian release.

Software Support cont.

The iso file can be burned to a cdrom that can be used to boot up a workstation.

Alternatively, virtualisation software such as

- ▶ *VMware Player / Workstation* (Windows, Linux),
- ▶ *VMware Fusion* (Mac OS X),
- ▶ *VirtualBox* (Windows, Linux) or
- ▶ *QEMU* (Linux)

can be used to run a 'virtual' guest computer alongside the host computer.

The software can also be installed to disk and updated using standard Debian tools; see the documentation for the 'Debian Live' tools used.

Appendix: Software Support cont.

Known issues with the provided iso file are:

- ▶ The cdrom appears to fail on some Dell models, there may be a BIOS incompatibility with the syslinux bootloader. Failures with the Parallels virtualisation for OS X were also reported, at least for the August 2008 release.
- ▶ No wireless extensions: if a laptop is booted off a cdrom, chances are that wireless will not be supported due to lack of binary firmware.

We will provide some examples or tips for Virtualbox, in particular the so-called Open Source Edition virtualbox-ose, that is readily available for Debian and Ubuntu.

Debian Live

The cdrom / iso was built using the excellent Debian Live toolkit.

Consequently, the documentation for Debian Live at <http://wiki.debian.org/DebianLive/> should be consulted in case of questions.

In particular, the FAQ at <http://wiki.debian.org/DebianLive/FAQ> is quite helpful.

Virtualbox networking

By default, virtualbox sets up networking using NAT: the virtualbox client can go 'outside', but cannot be seen from the outside due to the NAT layer. This is similar to how home networking is often invisible behind a single wired or wireless router with NAT capabilities.

When using Linux as a host operating system, it is possible to set up bridge networking where the virtualbox client becomes visible to the network just like another machine. See for example the short tutorial at

```
http://www.ubuntugeek.com/how-to-set-up-host-interface-networking-for-virtualbox-on-ubuntu.html.
```

Virtualbox networking

For example, on a Debian/Ubuntu host, the networking can be configured via

```
auto br0
    iface br0 inet static
        address 192.168.1.10
        netmask 255.255.255.0
        network 192.168.1.0
        broadcast 192.168.1.255
        gateway 192.168.1.4
        bridge_ports eth0 vbox0
```

which sets up a bridge interface to two interfaces `vbox0` and `vbox1`. These can then be assigned as host interface in the virtualbox console.

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

bigm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Profiling

We need to know where our code spends the time it takes to compute our tasks. Measuring is critical.

R already provides the basic tools for performance analysis.

- ▶ the `system.time` function for simple measurements.
- ▶ the `Rprof` function for profiling R code.
- ▶ the `Rprofmem` function for profiling R memory usage.

In addition, the `profr` package on CRAN can be used to visualize `Rprof` data.

Profiling cont.

The chapter *Tidying and profiling R code* in the *R Extensions* manual is a good first source for documentation on profiling and debugging.

Simon Urbanek has a page on benchmarks (for Macs) at <http://r.research.att.com/benchmarks/>

Lastly, we can also profile compiled code.

The following example (taken from the manual) contains two calls to `Rprof` to turn profiling on and off.

RProf example

```
library(MASS); library(boot)
storm.fm <- nls(Time ~ b*Viscosity/(Wt - c), stormer, \
               start = c(b=29.401, c=2.2183))
st <- cbind(stormer, fit=fitted(storm.fm))
storm.bf <- function(rs, i) {
  st$Time <- st$fit + rs[i]
  tmp <- nls(Time ~ (b * Viscosity)/(Wt - c), st, \
             start = coef(storm.fm))
  tmp$m$getAllPars()
}
# remove mean
rs <- scale(resid(storm.fm), scale = FALSE)
Rprof("boot.out")
# pretty slow
storm.boot <- boot(rs, storm.bf, R = 4999)
Rprof(NULL)
```


RProf example cont.

We can run the example via either one of

```
cat profilingExample.R | R --no-save # N = 4999
cat profilingSmall.R | R --no-save # N = 99
```

We can then analyse the output using two different ways. First, directly from R into an R object:

```
data <- summaryRprof("boot.out")
print(str(data))
```

Second, from the command-line (on systems having Perl)

```
R CMD Prof boot.out | less
```

Third, `profr` can directly profile, evaluate, and optionally plot, an expression. Note that we reduce N here:

```
pr <- profr(storm.boot <- boot(rs, storm.bf, R = 199))
plot(pr)
ggplot(pr) ## using library(ggplot2)
```

In this example, the code is already very efficient and no 'smoking gun' reveals itself for further improvement.

profr example

The `profr` function can be very useful for its quick visualisation of the `RProf` output. Consider this contrived example:

```
sillysum <- function(N) {s <- 0;
  for (i in 1:N) s <- s + i; s}
ival <- 1/5000
Rprof("/tmp/sillysum.out", interval=ival)
a <- sillysum(1e6); Rprof(NULL)
plot(parse_rprof("/tmp/sillysum.out", interval=ival))
```

and a more efficient solution where we use a larger N :


```
efficientsum <- function(N) { s <- sum(seq(1,N)); s }
ival <- 1/5000
Rprof("/tmp/effsum.out", interval=ival)
a <- efficientsum(1e7); Rprof(NULL)
plot(parse_rprof("/tmp/effsum.out", interval=ival))
```

We can run the complete example via


```
cat rprofChartExample.R | R --no-save
```

profr example cont.

Profile of inefficient summation



Profile of efficient summation



RProfmem example

When R has been built with the `enable-memory-profiling` option, we can also look at use of memory and allocation.

To continue with the *R Extensions* manual example, we issue calls to `Rprofmem` to start and stop logging to a file as we did for `Rprof`:

```
Rprofmem("/tmp/boot.memprof", threshold=1000)
storm.boot <- boot(rs, storm.bf, R = 4999)
Rprofmem(NULL)
```

Looking at the results files shows, and we quote, that *apart from some initial and final work in 'boot' there are no vector allocations over 1000 bytes.*

We also mention in passing that the `tracemem` function can log when copies of a (presumably large) object are being made. Details are in section 3.3.3 of the *R Extensions* manual.

Profiling compiled code

Profiling compiled code typically entails rebuilding the binary and libraries with the `-gp` compiler option. In the case of `R`, a complete rebuild is required.

Add-on tools like `valgrind` and `kcachegrind` can be helpful.

Two other options are mentioned in the *R Extensions* manual section of profiling for Linux.

First, `sprof`, part of the C library, can profile shared libraries. Second, the add-on package `oprofile` provides a daemon that has to be started (stopped) when profiling data collection is to start (end).

A third possibility is the use of the Google Perftools package which we will illustrate.

Profiling with Google Perftools

The Google Perftools package provides four modes of performance analysis / improvement:

- ▶ a thread-caching malloc (memory allocator),
- ▶ a heap-checking facility,
- ▶ a heap-profiling facility and
- ▶ cpu profiling.

Here, we will focus on the last feature.

There are two possible modes of running code with the cpu profile.

The preferred approach is to link with `-lprofiler`.

Alternatively, one can dynamically pre-load the profiler library.

Profiling with Google Perftools

```
# turn on profiling and provide a profile log file
LD_PRELOAD="/usr/lib/libprofiler.so.0" \
CPUPROFILE=/tmp/rprof.log \
r profilingSmall.R
```

We can then analyse the profiling output in the file. The profiler (renamed from `pprof` to `google-pprof` on Debian) has a large number of options. Here just use two different formats:

```
# show text output
google-pprof --cum --text \
  /usr/bin/r /tmp/rprof.log | less
```

```
# or analyse call graph using gv
google-pprof --gv /usr/bin/r /tmp/rprof.log
```

The shell script `googlePerftools.sh` runs the complete example.

Profiling with Google Perftools

Another output for format is for the *callgrind* analyser that is part of *valgrind*—a frontend to a variety of analysis tools such as *cachegrind* (cache simulator), *callgrind* (call graph tracer), *helpgrind* (race condition analyser), *massif* (heap profiler), and *memcheck* (fine-grained memory checker).

For example, the KDE frontend *kcachegrind* can be used to visualize the profiler output as follows:

```
google-pprof --callgrind \  
  /usr/bin/r /tmp/gpProfile.log \  
  > googlePerftools.callgrind  
kcachegrind googlePerftools.callgrind
```

Profiling with Google Perftools

Kcachegrind running on the the profiling output looks as follows:

googlePerftoolsGraphviz.callgrind - KCachegrind <@ron>

File View Go Settings Help

Hits

Search: (No Grouping)

| Incl. | Self | Called | Function |
|---------|------|--------|------------------------|
| 2555.56 | 7.41 | 769 | .L1191 |
| 748.15 | 0.00 | 202 | do_begin |
| 618.52 | 3.70 | 167 | Rf_applyClosure |
| 314.81 | 0.00 | 85 | forcePromise |
| 214.81 | 0.00 | 58 | do_set |
| 181.48 | 0.00 | 49 | .L1573 |
| 177.78 | 3.70 | 48 | do_internal |
| 125.93 | 0.00 | 34 | Rf_evalList |
| 77.78 | 0.00 | 21 | .L1560 |
| 66.67 | 0.00 | 18 | do_if |
| 55.56 | 0.00 | 15 | .L1508 |
| 40.74 | 0.00 | 11 | Rf_evalListKeepMissing |
| 37.04 | 0.00 | 10 | do_return |
| 18.52 | 7.41 | 6 | _init <cycle 1> |
| 18.52 | 7.41 | 2 | <cycle 1> |
| 18.52 | 3.70 | 5 | Rf_Scollate |
| 18.52 | 0.00 | 5 | .L939 |
| 18.52 | 0.00 | 5 | .L1667 |
| 18.52 | 0.00 | 5 | .L1609 |
| 18.52 | 0.00 | 5 | .L1200 |

.L1191

Types Callers All Callers Source Callee Map

Hits Assembler Source Positio

1 There is no instruction info in the profile data file.

Caller Map Parts Call Graph Callees All Callees Assembler

googlePerftoolsGraphviz.callgrind [1] - Total Hits Cost: 27

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Vectorisation

Revisiting our trivial trivial example from the preceding section:

```
> sillysum <- function(N) { s <- 0;
  for (i in 1:N) s <- s + i; return(s) }
> system.time(print(sillysum(1e7)))
```

```
[1] 5e+13
   user  system elapsed
13.617   0.020  13.701
```

```
>
```

```
> system.time(print(sum(as.numeric(seq(1,1e7)))))
```

```
[1] 5e+13
   user  system elapsed
 0.224   0.092   0.315
```

```
>
```

Replacing the loop yielded a gain of a factor of more than 40.
Hence it pays to know the corpus of available functions.

Vectorisation cont.

A more interesting example is provided in a [case study](#) on the [Ra](#) (c.f. next section) site and taken from the *S Programming* book:

Consider the problem of finding the distribution of the determinant of a 2×2 matrix where the entries are independent and uniformly distributed digits $0, 1, \dots, 9$. This amounts to finding all possible values of $ac - bd$ where a, b, c and d are digits.

Vectorisation cont.

The brute-force solution is using explicit loops over all combinations:

```
dd.for.c <- function() {  
  val <- NULL  
  for (a in 0:9) for (b in 0:9)  
    for (d in 0:9) for (e in 0:9)  
      val <- c(val, a*b - d*e)  
  table(val)  
}
```

The naive time is

```
> mean(replicate(10, system.time(dd.for.c())["elapsed"]))  
[1] 0.2678
```

Vectorisation cont.

The case study discusses two important points that bear repeating:

- ▶ pre-allocating space helps with performance:

```
val <- double(10000)
```

reduces the time to 0.1204

- ▶ switching to faster functions can help too as `tabulate` outperforms `table` and reduced the time further to 0.1180.

Vectorisation cont.

However, by far the largest improvement comes from eliminating the four loop with two calls each to `outer`:

```
dd.fast.tabulate <- function() {
  val <- outer(0:9, 0:9, "*")
  val <- outer(val, val, "-")
  tabulate(val)
}
```

The time for the most efficient solution is:

```
> mean(replicate(10,
  system.time(dd.fast.tabulate())["elapsed"]))
```

```
[1] 0.0014
```

which is orders of magnitude faster.

All examples can be run via the script `dd.naive.r`.

Accelerated R with just-in-time compilation

Stephen Milborrow recently released a set of patches to R that allow 'just-in-time compilation' of loops and arithmetic expression. Together with his `jit` package on CRAN, this can be used to obtain speedups of standard R operations.

Our trivial example run in Ra:

```
library(jit)
sillysum <- function(N) { jit(1); s <- 0; \
  for (i in 1:N) s <- s + i; return(s) }

> system.time(print(sillysum(1e7)))
[1] 5e+13
   user  system elapsed
 1.548   0.028   1.577
```

which gets a speed increase of a factor of five – not bad at all.

Accelerated R with just-in-time compilation

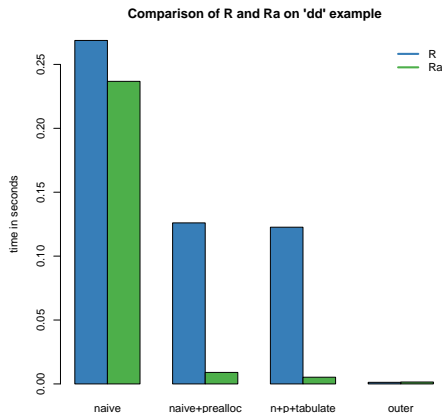
The last looping example can be improved with jit:

```
dd.for.pre.tabulate.jit <- function() {
  jit(1)
  val <- double(10000)
  i <- 0
  for (a in 0:9) for (b in 0:9)
    for (d in 0:9) for (e in 0:9) {
      val[i <- i + 1] <- a*b - d*e
    }
  tabulate(val)
}
```

```
> mean(replicate(10, system.time(dd.for.pre.tabulate.jit()))["
[1] 0.0053
```

or only about three to four times slower than the non-looped solution using 'outer'—a rather decent improvement.

Accelerated R with just-in-time compilation



Ra achieves very good decreases in total computing time in these examples but cannot improve the efficient solution any further.

Ra and `jit` are still fairly new and not widely deployed yet, but readily available in Debian and Ubuntu.

Source: Our calculations

Optimised Blas

Blas ('basic linear algebra subprogram', see [Wikipedia](#)) are standard building block for linear algebra. Highly-optimised libraries exist that can provide considerable performance gains.

[R](#) can be built using so-called optimised Blas such as Atlas ('free'), Goto (not 'free'), or those from Intel or AMD; see the 'R Admin' manual, section A.3 'Linear Algebra'.

The speed gains can be noticeable. For Debian/Ubuntu, one can simply install one of the `atlas-base-*` packages.

An example from the old README.Atlas, running with R 2.7.0, follows:

Optimised Blas cont.

```
# with Atlas
> mm <- matrix(rnorm(4*10^6), ncol = 2*10^3)
> mean(replicate(10,
  system.time(crossprod(mm)) ["elapsed"]))
```

```
[1] 3.8465
```

```
# with basic. non-optimised Blas,
# ie after dpkg --purge atlas3-base libatlas3gf-base
> mm <- matrix(rnorm(4*10^6), ncol = 2*10^3)
> mean(replicate(10,
  system.time(crossprod(mm)) ["elapsed"]))
```

```
[1] 8.9776
```

So for pure linear algebra problems, we may get an improvement by a factor of two or larger by using binary code that is optimised for the cpu class. This is likely to be more pronounced on multi-cpu machines.

Higher increases are possibly by 'tuning' the library, see the Atlas documentation.

From Blas to GPUs.

The next frontier for hardware acceleration is computing on GPUs ('graphics programming units', see [Wikipedia](#)).

GPUs are essentially hardware that is optimised for both I/O and floating point operations, leading to much faster code execution than standard CPUs on floating-point operations.

Development kits are available (e.g Nvidia CUDA) and the recently announced OpenCL programming specification should make GPU-computing vendor-independent.

Some initial work on integration with [R](#) has been undertaken but there appear to no easy-to-install and easy-to-use kits for [R](#) – yet.

So this provides a perfect intro for the next subsection on compilation.

Compiled Code

Beyond smarter code (using e.g. vectorised expression and/or just-in-time compilation), compiled subroutines or accelerated libraries, the most direct speed gain is to switch to compiled code.

This section covers two possible approaches:

- ▶ `inline` for automated wrapping of simple expression
- ▶ `Rcpp` for easing the interface between `R` and `C++`

Another different approach is to keep the core logic 'outside' but to *embed* `R` into the application. There is some documentation in the 'R Extensions' manual, and packages like `RApache` or `littler` offer concrete examples. This does however require a greater familiarity with `R` internals.

Compiled Code: The Basics

R offers several functions to access compiled code: `.C` and `.Fortran` as well as `.Call` and `.External`. (*R Extensions*, sections 5.2 and 5.9; *Software for Data Analysis*). `.C` and `.Fortran` are older and simpler, but more restrictive in the long run.

The canonical example in the documentation is the convolution function:

```
1 void convolve(double *a, int *na, double *b,  
2             int *nb, double *ab)  
3 {  
4     int i, j, nab = *na + *nb - 1;  
5  
6     for(i = 0; i < nab; i++)  
7         ab[i] = 0.0;  
8     for(i = 0; i < *na; i++)  
9         for(j = 0; j < *nb; j++)  
10            ab[i + j] += a[i] * b[j];  
11 }
```


Compiled Code: The Basics cont.

The convolution function is called from R by

```
1 conv <- function(a, b)
2   .C("convolve",
3     as.double(a),
4     as.integer(length(a)),
5     as.double(b),
6     as.integer(length(b)),
7     ab = double(length(a) + length(b) - 1))$ab
```

As stated in the manual, one must take care to coerce all the arguments to the correct R storage mode before calling `.C` as mistakes in matching the types can lead to wrong results or hard-to-catch errors.

The script `convolve.C.sh` compiles and links the source code, and then calls R to run the example.

Compiled Code: The Basics cont.

Using `.Call`, the example becomes

```
1 #include <R.h>
2 #include <Rdefines.h>
3
4 SEXP convolve2(SEXP a, SEXP b)
5 {
6     int i, j, na, nb, nab;
7     double *xa, *xb, *xab;
8     SEXP ab;
9
10    PROTECT(a = AS_NUMERIC(a));
11    PROTECT(b = AS_NUMERIC(b));
12    na = LENGTH(a); nb = LENGTH(b); nab = na + nb - 1;
13    PROTECT(ab = NEW_NUMERIC(nab));
14    xa = NUMERIC_POINTER(a); xb = NUMERIC_POINTER(b);
15    xab = NUMERIC_POINTER(ab);
16    for(i = 0; i < nab; i++) xab[i] = 0.0;
17    for(i = 0; i < na; i++)
18        for(j = 0; j < nb; j++) xab[i + j] += xa[i] * xb[j];
19    UNPROTECT(3);
20    return(ab);
21 }
```



Compiled Code: The Basics cont.

Now the call simply becomes easier using the function name and the vector arguments as all handling is done at the C/C++ level:

```
conv <- function(a, b) .Call("convolve2", a, b)
```

The script `convolve.Call.sh` compiles and links the source code, and then calls **R** to run the example.

In summary, we see that

- ▶ there are different entry points
- ▶ using different calling conventions
- ▶ leading to code that may need to do more work at the lower level.

Compiled Code: inline

`inline` is a package by Oleg Sklyar et al that provides the function `cfunction` that can wrap Fortran, C or C++ code.

```
1 ## A simple Fortran example
2 code <- "
3     integer i
4     do 1 i=1, n(1)
5     1 x(i) = x(i)**3
6 "
7 cubefn <- cfunction(signature(n="integer", x="numeric"),
8                     code, convention=".Fortran")
9 x <- as.numeric(1:10)
10 n <- as.integer(10)
11 cubefn(n, x)$x
```

`cfunction` takes care of compiling, linking, loading, ... by placing the resulting dynamically-loadable object code in the per-session temporary directory used by R.

Run this via `cat inline.Fortan.R | R -no-save`.

Compiled Code: inline cont.

`inline` defaults to using the `.Call()` interface:

```

1  ## Use of .Call convention with C code
2  ## Multiplying each image in a stack with a 2D Gaussian at a given position
3  code <- "
4    SEXP res;
5    int nprotect = 0, nx, ny, nz, x, y;
6    PROTECT(res = Rf_duplicate(a)); nprotect++;
7    nx = INTEGER(GET_DIM(a))[0];
8    ny = INTEGER(GET_DIM(a))[1];
9    nz = INTEGER(GET_DIM(a))[2];
10   double sigma2 = REAL(s)[0] * REAL(s)[0], d2 ;
11   double cx = REAL(centre)[0], cy = REAL(centre)[1], *data, *rdata;
12   for (int im = 0; im < nz; im++) {
13     data = &(REAL(a)[im*nx*ny]); rdata = &(REAL(res)[im*nx*ny]);
14     for (x = 0; x < nx; x++)
15       for (y = 0; y < ny; y++) {
16         d2 = (x-cx)*(x-cx) + (y-cy)*(y-cy);
17         rdata[x + y*nx] = data[x + y*nx] * exp(-d2/sigma2);
18       }
19   }
20   UNPROTECT(nprotect);
21   return res;
22 "
23 funx <- cfunction(signature(a="array", s="numeric", centre="numeric"), code)
24
25 x <- array(runif(50*50), c(50,50,1))
26 res <- funx(a=x, s=10, centre=c(25,15))  ## actual call of compiled function
27 if (interactive()) image(res[, ,1])

```

Compiled Code: inline cont.

We can revisit the earlier distribution of determinants example. If we keep it very simple and pre-allocate the temporary vector in `R`, the example becomes

```
1 code <- "  
2   if (isNumeric(vec)) {  
3     int *pv = INTEGER(vec);  
4     int n = length(vec);  
5     if (n = 10000) {  
6       int i = 0;  
7       for (int a = 0; a < 9; a++)  
8         for (int b = 0; b < 9; b++)  
9           for (int c = 0; c < 9; c++)  
10            for (int d = 0; d < 9; d++)  
11              pv[i++] = a*b - c*d;  
12     }  
13   }  
14   return(vec);  
15 "  
16  
17 funx <- cfunction(signature(vec="numeric"), code)
```

Compiled Code: inline cont.

We can use the inlined function in new function to be timed:

```
dd.inline <- function() {  
  x <- integer(10000)  
  res <- funx(vec=x)  
  tabulate(res)  
}  
> mean(replicate(100, system.time(dd.inline()))["elapsed"]  
[1] 0.00051
```

Even though it uses the simplest algorithm, pre-allocates memory in R and analyse the result in R, it still more than twice as fast than the previous best solution.

The script `dd.inline.r` runs this example.

Compiled Code: Rcpp

[Rcpp](#) makes it easier to interface C++ and [R](#) code.

Using the `.Call` interface, we can use features of the C++ language to automate the tedious bits of the macro-based C-level interface to [R](#).

One major advantage of using `.Call` is that vectors (or matrices) can be passed directly between [R](#) and C++ without the need for explicit passing of dimension arguments. And by using the C++ class layers, we do not need to directly manipulate the SEXP objects.

So let us rewrite the 'distribution of determinant' example one more time.

Rcpp example

The simplest version can be set up as follows:

```
1 #include <Rcpp.hpp>
2
3 RcppExport SEXP dd_rcpp(SEXP v) {
4   SEXP r1 = R_NilValue;          // Use this when there is nothing to be returned
5   .
6   RcppVector<int> vec(v);        // vec parameter viewed as vector of doubles.
7   int n = vec.size(), i = 0;
8
9   for (int a = 0; a < 9; a++)
10    for (int b = 0; b < 9; b++)
11     for (int c = 0; c < 9; c++)
12      for (int d = 0; d < 9; d++)
13        vec(i++) = a*b - c*d;
14
15   RcppResultSet rs;             // Build result set to be returned as a list to
16   R.                             // R.
17   rs.add("vec", vec);           // vec as named element with name 'vec'
18   r1 = rs.getReturnList();      // Get the list to be returned to R.
19
20   return r1;
}
```

but it is actually preferable to use the exception-handling feature of C++ as in the slightly longer next version.

Rcpp example cont.

```
1 #include <Rcpp.hpp>
2
3 RcppExport SEXP dd_rcpp(SEXP v) {
4   SEXP r1 = R_NilValue; // Use this when there is nothing to be returned.
5   char* exceptionMesg = NULL; // msg var in case of error
6
7   try {
8     RcppVector<int> vec(v); // vec parameter viewed as vector of doubles.
9     int n = vec.size(), i = 0;
10    for (int a = 0; a < 9; a++)
11      for (int b = 0; b < 9; b++)
12        for (int c = 0; c < 9; c++)
13          for (int d = 0; d < 9; d++)
14            vec(i++) = a*b - c*d;
15
16    RcppResultSet rs; // Build result set to be returned as a list to R.
17    rs.add("vec", vec); // vec as named element with name 'vec'
18    r1 = rs.getReturnList(); // Get the list to be returned to R.
19  } catch (std::exception& ex) {
20    exceptionMesg = copyMessageToR(ex.what());
21  } catch (...) {
22    exceptionMesg = copyMessageToR("unknown reason");
23  }
24
25  if (exceptionMesg != NULL)
26    error(exceptionMesg);
27
28  return r1;
29 }
```

Rcpp example cont.

We can create a shared library from the source file as follows:

```
PKG_CPPFLAGS='r -e' cat (Rcpp:::RcppCxxFlags()) ' ' \  
  R CMD SHLIB dd.rcpp.cpp \  
  `r -e' cat (Rcpp:::RcppLdFlags()) ' ' \  
  
g++ -I/usr/share/R/include \  
  -I/usr/lib/R/site-library/Rcpp/lib \  
  -fpic -g -O2 \  
  -c dd.rcpp.cpp -o dd.rcpp.o \  
g++ -shared -o dd.rcpp.so dd.rcpp.o \  
  -L/usr/lib/R/site-library/Rcpp/lib \  
  -lRcpp -Wl,-rpath,/usr/lib/R/site-library/Rcpp/lib \  
  -L/usr/lib/R/lib -lR
```

Note how we let the [Rcpp](#) package tell us where header and library files are stored.

Rcpp example cont.

We can then load the file using `dyn.load` and proceed as in the `inline` example.

```
dyn.load("dd.rcpp.so")
```

```
dd.rcpp <- function() {  
  x <- integer(10000)  
  res <- .Call("dd_rcpp", x)  
  tabulate(res$vec)  
}
```

```
mean(replicate(100, system.time(dd.rcpp())["elapsed"]))  
[1] 0.00047
```

This beats the `inline` example by a negligible amount which is probably due to some overhead in the easy-to-use inlining.

The file `dd.rcpp.sh` runs the full Rcpp example.

Rcpp example cont.

Two tips for easing builds with `Rcpp`:

For command-line use, it is easiest to link `Rcpp.h` to `/usr/local/include`, and `libRcpp.so` to `/usr/local/lib`. The example reduces to

```
R CMD SHLIB dd.rcpp.cpp
```

as header and library will be found in the default locations.

For package building, we can have a file `src/Makevars` with

```
# compile flag providing header directory
PKG_CXXFLAGS='Rscript -e 'cat(Rcpp:::RcppCxxFlags())'`
# link flag providing library and path
PKG_LIBS='Rscript -e 'cat(Rcpp:::RcppLdFlags())'`
```

See the `help(Rcpp-package)` page for more.

Debugging example: valgrind

Analysis of compiled code is mainly undertaken with a debugger like [gdb](#), or a graphical frontend like [ddd](#).

Another useful tool is [valgrind](#) which can find memory leaks. We can illustrate its use with a recent real-life example.

[RMySQL](#) had recently been found to be leaking memory when database connections are being established and closed. Given how [RPostgreSQL](#) shares a common heritage, it seemed like a good idea to check.

Debugging example: valgrind

We create a small test script which opens and closes a connection to the database in a loop and sends a small 'select' query. We can run this in a way that is close to the suggested use from the 'R Extensions' manual:

```
R -d "valgrind -tool=memcheck  
-leak-check=full" -vanilla < valgrindTest.R
```

which creates copious output, including what is on the next slide.

Given the source file and line number, it is fairly straightforward to locate the source of error: a vector of pointers was freed without freeing the individual entries first.

Debugging example: valgrind

The state before the fix:

```
[...]  
#==21642== 2,991 bytes in 299 blocks are definitely lost in loss record 34 of 47  
#==21642==   at 0x4023D6E: malloc (vg_replace_malloc.c:207)  
#==21642==   by 0x6781CAF: RS_DBI_copyString (RS-DBI.c:592)  
#==21642==   by 0x6784B91: RS_PostgreSQL_createDataMappings (RS-PostgreSQL.c:400)  
#==21642==   by 0x6785191: RS_PostgreSQL_exec (RS-PostgreSQL.c:366)  
#==21642==   by 0x40C50BB: (within /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40EDD49: Rf_eval (in /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40F00DC: (within /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40EDA74: Rf_eval (in /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40F0186: (within /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40EDA74: Rf_eval (in /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40F16E6: Rf_applyClosure (in /usr/lib/R/lib/libR.so)  
#==21642==   by 0x40ED99A: Rf_eval (in /usr/lib/R/lib/libR.so)  
#==21642==  
#==21642== LEAK SUMMARY:  
#==21642==   definitely lost: 3,063 bytes in 301 blocks.  
#==21642==   indirectly lost: 240 bytes in 20 blocks.  
#==21642==   possibly lost: 9 bytes in 1 blocks.  
#==21642==   still reachable: 13,800,378 bytes in 8,420 blocks.  
#==21642==   suppressed: 0 bytes in 0 blocks.  
#==21642== Reachable blocks (those to which a pointer was found) are not shown.  
#==21642== To see them, rerun with: --leak-check=full --show-reachable=yes
```


Debugging example: valgrind

The state after the fix:

```
[...]
===3820===
===3820=== 312 (72 direct, 240 indirect) bytes in 2 blocks are definitely lost in loss record 1
===3820===   at 0x4023D6E: malloc (vg_replace_malloc.c:207)
===3820===   by 0x43F1563: nss_parse_service_list (nsswitch.c:530)
===3820===   by 0x43F1CC3: __nss_database_lookup (nsswitch.c:134)
===3820===   by 0x445EF4B: ???
===3820===   by 0x445FCEC: ???
===3820===   by 0x43AB0F1: getpwuid_r@@GLIBC_2.1.2 (getXXbyYY_r.c:226)
===3820===   by 0x43AAA76: getpwuid (getXXbyYY.c:116)
===3820===   by 0x4149412: (within /usr/lib/R/lib/libR.so)
===3820===   by 0x412779D: (within /usr/lib/R/lib/libR.so)
===3820===   by 0x40EDA74: Rf_eval (in /usr/lib/R/lib/libR.so)
===3820===   by 0x40F00DC: (within /usr/lib/R/lib/libR.so)
===3820===   by 0x40EDA74: Rf_eval (in /usr/lib/R/lib/libR.so)
===3820===
===3820=== LEAK SUMMARY:
===3820===   definitely lost: 72 bytes in 2 blocks.
===3820===   indirectly lost: 240 bytes in 20 blocks.
===3820===   possibly lost: 0 bytes in 0 blocks.
===3820===   still reachable: 13,800,378 bytes in 8,420 blocks.
===3820===   suppressed: 0 bytes in 0 blocks.
===3820=== Reachable blocks (those to which a pointer was found) are not shown.
===3820=== To see them, rerun with: --leak-check=full --show-reachable=yes
```

showing that we recovered 3000 bytes.

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Embarassingly parallel

Several R packages on CRAN provide the ability to execute code in parallel:

- ▶ NWS
- ▶ Rmpi
- ▶ snow (using MPI, PVM, NWS or sockets)
- ▶ papply
- ▶ taskPR

A recent paper (Schmidberger, Morgan, Eddelbuettel, Yu, Rossini, Tierney and Mansmann, 2008) surveys the field.

NWS Intro

NWS, or NetWorkSpaces, is an alternative to MPI (which we discuss below). Based on Python, it may be easier to install (in case administrator rights are unavailable) and use than MPI. It is accessible from R, Python and Matlab. It is also cross-platform.

NWS is available via [Sourceforge](#) as well as [CRAN](#). An introductory article (focussing on Python) appeared last summer in [Dr. Dobb's](#).

On Debian and Ubuntu, installing the `python-nwserver` package on at least the server node, and installing `r-cran-nws` on each client is all that is needed. Other system may need to install the `twisted` framework for `Python` first.

NWS data store example

A simple example, adapted from one of the package demos:

```
ws <- netWorkspace('r place') # create a 'value store'
nwsStore(ws, 'x', 1)           # place a value (as fifo)

cat(nwsListVars(ws), "\n")    # we can list
nwsFind(ws, 'x')              # and lookup
nwsStore(ws, 'x', 2)          # and overwrite
cat(nwsListVars(ws), "\n")    # now see two entries

cat(nwsFetch(ws, 'x'), '\n')  # we can fetch
cat(nwsFetch(ws, 'x'), '\n')  # we can fetch
cat(nwsListVars(ws), '\n')    # and none left

cat(nwsFetchTry(ws, 'x', 'no go'), '\n') # can't fetch
```

The script `nwsVariableStore.r` contains this and a few more commands.

NWS sleigh example

The NWS component sleigh is an R class that makes it very easy to write simple parallel programs. Sleigh uses the master/worker paradigm: The master submits tasks to the workers, who may or may not be on the same machine as the master.

```
# create a sleigh object on two nodes using ssh
s <- sleigh(nodeList=c("joe", "ron"), launch=sshcmd)

# execute a statement on each worker node
eachWorker(s, function() x <<- 1)

# get system info from each worker
eachWorker(s, Sys.info)

# run a lapply-style funct. over each list elem.
eachElem(s, function(x) {x+1}, list(1:10))

stopSleigh(s)
```

NWS sleigh cont.

Also of note is the extended `caretNWS` version of `caret` by Max Kuhn, and described in a recent Journal of Statistical Software article.

`caret` (short for 'Classification and Regression Training') provides a consistent interface for dozens of modern regression and classification techniques.

`caretNWS` uses `nws` and `sleigh` to execute embarrassingly parallel tasks: bagging, boosting, cross-validation, bootstrapping, ... This is all done 'behind the scenes' and thus easy to deploy.

Schmidberger et al find NWS to be competitive with the other parallel methods for non-degenerate cases where the ratio between communication and computation is balanced.

Rmpi

[Rmpi](#) is a CRAN package that provides an interface between [R](#) and the Message Passing Interface (MPI), a [standard](#) for parallel computing. (c.f. [Wikipedia](#) for more and links to the Open MPI and MPICH2 projects for implementations).

The preferred implementation for MPI is now [Open MPI](#). However, the older LAM implementation can be used on those platforms where Open MPI is unavailable. There is also an alternate implementation called MPICH2. Lastly, we should also mention the similar Parallel Virtual Machine (PVM) tool; see its [Wikipedia](#) page for more.

[Rmpi](#) allows us to use MPI directly from [R](#) and comes with several examples. However, we will focus on the higher-level usage via [snow](#).

MPI Example

Let us look at the `MPI` variant of the 'Hello, World!' program:

```
1 #include <stdio.h>
2 #include "mpi.h"
3
4 int main(int argc, char** argv)
5 {
6     int rank, size, nameLen;
7     char processorName[MPI_MAX_PROCESSOR_NAME];
8
9     MPI_Init(&argc, &argv);
10    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
11    MPI_Comm_size(MPI_COMM_WORLD, &size);
12
13    MPI_Get_processor_name(processorName, &nameLen);
14
15    printf("Hello, rank %d, size %d on processor %s\n",
16          rank, size, processorName);
17
18    MPI_Finalize();
19    return 0;
20 }
```

MPI Example: cont.

We can compile the previous example via

```
$ mpicc -o mpiHelloWorld mpiHelloWorld.c
```

If it has been copied across several Open MPI-equipped hosts, we can execute it N times on the M listed hosts via:

```
$ orterun -n 8 -H ron,joe,wayne,tony /tmp/mpiHelloWorld
```

```
Hello, rank 0, size 8 on processor ron  
Hello, rank 4, size 8 on processor ron  
Hello, rank 6, size 8 on processor wayne  
Hello, rank 3, size 8 on processor tony  
Hello, rank 2, size 8 on processor wayne  
Hello, rank 5, size 8 on processor joe  
Hello, rank 7, size 8 on processor tony  
Hello, rank 1, size 8 on processor joe
```

Notice how the order of execution is indeterminate.

MPI Example: cont.

Besides `orterun` (which replaces the `mpirun` command used by other MPI implementations), Open MPI also supplies `ompi_info` to query parameter settings.

Open MPI has very fine-grained configuration options that permit e.g. attaching particular jobs to particular cpus or cores.

Detailed documentation is provided at the web site <http://www.openmpi.org>.

We will concentrate on using MPI via the `Rmpi` package.

Rmpi

Rmpi, a CRAN package by Hao Yu, wraps many of the MPI API calls for use by **R**.

The preceding example can be rewritten in **R** as

```
1 #!/usr/bin/env r
2
3 library(Rmpi) # calls MPI_Init
4
5 rk <- mpi.comm.rank(0)
6 sz <- mpi.comm.size(0)
7 name <- mpi.get.processor.name()
8 cat("Hello , rank", rk, "size", sz, "on", name, "\n")
```

Rmpi: cont.

```
$ orterun -n 8 -H ron,joe,wayne,tony \  
/tmp/mpiHelloWorld.r
```

```
Hello, rank 0 size 8 on ron  
Hello, rank 4 size 8 on ron  
Hello, rank 3 size 8 on tony  
Hello, rank 7 size 8 on tony  
Hello, rank 6 size 8 on wayne  
Hello, rank 2 size 8 on wayne  
Hello, rank 5 size 8 on joe  
Hello, rank 1 size 8 on joe
```

Rmpi: cont.

We can also execute this as a one-liner using `r` (which we discuss later):

```
$ orterun -n 8 -H ron,joe,wayne,tony\  
  r -lRmpi -e'cat("Hello", \  
    mpi.comm.rank(0), "of", \  
    mpi.comm.size(0), "on", \  
    mpi.get.processor.name(), "\n")'
```

```
Hello, rank 0 size 8 on ron  
Hello, rank 4 size 8 on ron  
Hello, rank 3 size 8 on tony  
Hello, rank 7 size 8 on tony  
Hello, rank 6 size 8 on wayne  
Hello, rank 2 size 8 on wayne  
Hello, rank 5 size 8 on joe  
Hello, rank 1 size 8 on joe
```

Rmpi: cont.

`Rmpi` offers a large number functions, mirroring the rich API provided by MPI.

`Rmpi` also offers extensions specific to working with `R` and its objects, including a set of `apply`-style functions to spread load across the worker nodes.

However, we will use `Rmpi` mostly indirectly via `snow`.

snow

The `snow` package by Tierney et al provides a convenient abstraction directly from R.

It can be used to initialize and use a compute cluster using one of the available methods direct socket connections, MPI, PVM, or (since the most recent release), NWS. We will focus on MPI.

A simple example:

```
cl <- makeCluster(4, "MPI")
print(clusterCall(cl, function() \
  Sys.info()[c("nodename", "machine")]))
stopCluster(cl)
```

which we can as a one-liner as shown on the next slide.

snow: Example

```
$ orterun -n 1 -H ron,joe r -lsnow,Rmpi \  
-e'cl <- makeCluster(4, "MPI"); \  
  res <- clusterCall(cl, \  
    function() Sys.info()["nodename"]); \  
  print(do.call(rbind,res)); \  
  stopCluster(cl)'
```

```
4 slaves are spawned successfully. 0 failed.
```

```
nodename
```

```
[1,] "joe"
```

```
[2,] "ron"
```

```
[3,] "joe"
```

```
[4,] "ron"
```

Note that we told `orterun` to start on only one node – as `snow` then starts four instances (which are split evenly over the two given hosts).

snow: Example cont.

The power of `snow` lies in the ability to use the `apply`-style paradigm over a cluster of machines:

```
params <- c("A", "B", "C", "D", "E", "F", "G", "H")
cl <- makeCluster(4, "MPI")
res <- parSapply(cl, params, \
                FUN=function(x) myBigFunction(x))
```

will 'unroll' the parameters `params` one-each over the function argument given, utilising the cluster `cl`. In other words, we will be running four copies of `myBigFunction()` at once.

So the `snow` package provides a unifying framework for parallelly executed `apply` functions.

We will come back to more examples with `snow` below.

papply, biopara and taskPR

We saw that `Rmpi` and `NWS` have `apply`-style functions, and that `snow` provides a unified layer. `papply` is another CRAN package that wraps around `Rmpi` to distribute processing `apply`-style functions across a cluster.

However, using the Open MPI-based `Rmpi` package, I was not able to get `papply` to actually successfully distribute – and retrieve – results across a cluster. So `snow` remains the preferred wrapper.

`biopara` is another package to distribute load across a cluster using direct socket-based communication. We consider `snow` to be a more general-purpose package for the same task.

`taskPR` uses the `MPI` protocol directly rather than via `Rmpi`. It is however hard-wired to use LAM and failed to launch under the Open MPI-implementation.

slurm resource management and queue system

Once the number of compute nodes increases, it becomes important to be able to allocate and manage resources, and to queue and batch jobs. A suitable tool is `slurm`, an open-source resource manager for Linux clusters.

Paraphrasing from the [slurm website](#):

- ▶ it allocates exclusive and/or non-exclusive access to resources (computer nodes) to users;
- ▶ it provides a framework for starting, executing, and monitoring (typically parallel) work on a set of allocated nodes.
- ▶ it arbitrates contention for resources by managing a queue of pending work.

Slurm is being developed by a consortium including LLNL, HP, Bull, and Linux Networks.

slurm example

Slurm is rather rich in features; we will only scratch the surface here.

Slurm can use many underlying message passing / communications protocols, and MPI is well supported.

In particular, Open MPI works well with slurm. That is an advantage inasmuch as it permits use of [Rmpi](#).

slurm example

A simple example:

```
$ srun -N 2 r -lRmpi -e'cat("Hello", \  
    mpi.comm.rank(0), "of", \  
    mpi.comm.size(0), "on", \  
    mpi.get.processor.name(), "\n")'
```

```
Hello 0 of 1 on ron
```

```
Hello 0 of 1 on joe
```

```
$ srun -n 4 -N 2 -O r -lRmpi -e'cat("Hello", \  
    mpi.comm.rank(0), "of", \  
    mpi.comm.size(0), "on", \  
    mpi.get.processor.name(), "\n")'
```

```
Hello 0 of 1 on ron
```

```
Hello 0 of 1 on ron
```

```
Hello 0 of 1 on joe
```

```
Hello 0 of 1 on joe
```

This shows how to *overcommit* jobs per node, and provides an example where we set the number of worker instances on the command-line.

slurm example

Additional command-line tools of interest are `salloc`, `sbatch`, `scontrol`, `squeue`, `scancel` and `sinfo`. For example, to see the status of a compute cluster:

```
$ sinfo
```

```
PARTITION AVAIL  TIMELIMIT NODES  STATE NODELIST
debug*      up    infinite     2   idle mccoymccoy,ron
```

This shows two idle nodes in a partition with the default name 'debug'.

The `sview` graphical user interface combines the functionality of a few of the command-line tools.

A more complete example will be provided below.

Using all those cores

Multi-core hardware is now a default, and the number of cores per cpus will only increase. It is therefore becoming more important for software to take advantage of these features.

Two recent (and still 'experimental') packages by Luke Tierney are addressing this question:

- ▶ `pnmath` uses OpenMP compiler directives for explicitly parallel code;
- ▶ `pnmath0` uses `pthread`s and implements the same interface.

They can be found at `http:`

`//www.stat.uiowa.edu/~luke/R/experimental/`

pnmath and pnmath0

Both `pnmath` and `pnmath0` provide parallelized vector math functions and support routines.

Upon loading either package, a number of vector math functions are replaced with versions that are parallelized using OpenMP. The functions will be run using multiple threads if their results will be long enough for the parallel overhead to be outweighed by the parallel gains. On load a calibration calculation is carried out to assess the parallel overhead and adjust these thresholds.

Profiling is probably the best way to assess the possible usefulness. As a quick illustration, we compute the `qtukey` function on a eight-core machine:

pnmath and pnmath0 illustration

```
$ r -e'N=1e3;print(system.time(qtukey(seq(1,N)/N,2,2)))'
```

```
   user  system elapsed
66.590   0.000  66.649
```

```
$ r -lpnmath -e'N=1e3; \
print(system.time(qtukey(seq(1,N)/N,2,2)))'
```

```
   user  system elapsed
67.580   0.080   9.938
```

```
$ r -lpnmath0 -e'N=1e3; \
print(system.time(qtukey(seq(1,N)/N,2,2)))'
```

```
   user  system elapsed
68.230   0.010   9.983
```

The 6.7-fold reduction in 'elapsed' time shows that the multithreaded version takes advantage of the 8 available cores at a sub-linear fashion as some communications overhead is involved.

These improvements will likely be folded into future R versions.

Scripting example for R and slurm

Being able to launch numerous R jobs in a parallel environments is helped by the ability to 'script' R.

Several simple methods existed to start R:

- ▶ `R CMD BATCH file.R`
- ▶ `echo "commands" | R -no-save`
- ▶ `R -no-save < file.R > file.Rout`

These are suitable for one-off scripts, but may be too fragile for distributed computing.

Use scripting with `r`

The `r` command of the `littler` package (as well as R's `Rscript`) provide more robust alternatives.

`r` can also be used four different ways:

- ▶ `r file.R`
- ▶ `echo "commands" | r`
- ▶ `r -lRmpi -e 'cat("Hello",
mpi.get.processor.name())'`
- ▶ and *shebang*-style in script files: `#!/usr/bin/r`

It is the last point that is of particular interest in this HPC context.

Also of note is the availability of the `getopt` package on CRAN.

slurm and snow

Having introduced `snow`, `slurm` and `r`, we would like to combine them.

However, there is are problems:

- ▶ `snow` has a master/worker paradigm yet `slurm` launches its nodes symmetrically,
- ▶ `slurm`'s `srun` has limits in spawning jobs
- ▶ with `srun`, we cannot communicate the number of nodes 'dynamically' into the script: `snow`'s cluster creation needs a hardwired number of nodes

slurm and snow solution

`snow` solves the master / worker problem by auto-discovery upon startup. The package contains two internal files `RMPISNOW` and `RMPISNOWprofile` that use a combination of shell and R code to determine the node identity allowing it to switch to master or worker functionality.

We can reduce the same problem to this for our R script:

```
ndsvpid <- Sys.getenv("OMPI_MCA_ns_nds_vpid")
if (ndsvpid == "0") {                                # are we the master ?
  makeMPIcluster()
} else {                                             # or are we a slave ?
  sink(file="/dev/null")
  slaveLoop(makeMPImaster())
  q()
}
```

slurm and snow solution

For example

```
1 #!/usr/bin/env r
2
3 suppressMessages(library(Rmpi))
4 suppressMessages(library(snow))
5
6 mpirank <- mpi.comm.rank(0) # just FYI
7 ndsvpid <- Sys.getenv("OMPI_MCA_ns_nds_vpid")
8 if (ndsvpid == "0") { # are we master ?
9   cat("Launching master", ndsvpid, "\n")
10  cl <- makeMPIcluster()
11 } else { # or are we a slave ?
12   cat("Launching slave", ndsvpid, "\n")
13   sink(file="/dev/null")
14   slaveLoop(makeMPImaster())
15   q()
16 }
17 stopCluster(cl)
```

slurm and snow solution

The example creates

```
$ orterun -H ron,joe -n 4 rMPIsnowSimple.r
```

```
Launching slave 2
```

```
Launching master 0
```

```
Launching slave 1
```

```
Launching slave 3
```

and we see that $N - 1$ workers are running with one instance running as the coordinating manager node.

salloc for snow

The other important aspect is to switch to `salloc` (which will call `orterun`) instead of `srun`.

We can either supply the hosts used using the `-w` switch, or rely on the `slurm.conf` file.

But importantly, we can govern from the call how many instances we want running (and have neither the `srun` limitation requiring overcommitting nor the hard-coded `snow` cluster-creation size):

```
$ salloc -w ron,mccoy orterun -n 7 rMPISnow.r
```

We ask for a `slurm` allocation on the given hosts, and instruct Open MPI to run seven instances.

salloc for snow

```
1 #!/usr/bin/env r
2
3 suppressMessages(library(Rmpi))
4 suppressMessages(library(snow))
5
6 ndsvpid <- Sys.getenv("OMPI_MCA_ns_nds_vpid")
7 if (ndsvpid == "0") { # are we master ?
8   makeMPIcluster()
9 } else { # or are we a slave ?
10   sink(file="/dev/null")
11   slaveLoop(makeMPImaster())
12   q()
13 }
14
15 ## a trivial main body
16 cl <- getMPIcluster()
17 clusterEvalQ(cl, options("digits.secs"=3)) # show millisec
18 res <- clusterCall(cl, function() \
19   paste(Sys.info()[ "nodename" ], format(Sys.time())))
20 print(do.call(rbind, res))
21 stopCluster(cl)
```

salloc for snow

```
$ salloc -w ron,joe -n 7 rMPISnow.r
```

```
salloc: Granted job allocation 8
```

```
[,1]
```

```
[1,] "joe 2008-12-19 21:18:08.787"
```

```
[2,] "ron 2008-12-19 21:18:08.997"
```

```
[3,] "joe 2008-12-19 21:18:08.771"
```

```
[4,] "ron 2008-12-19 21:18:09.011"
```

```
[5,] "joe 2008-12-19 21:18:08.781"
```

```
[6,] "ron 2008-12-19 21:18:09.014"
```

```
sall: Relinquishing job allocation 8
```

A complete example

```
cl <- NULL
ndsvpid <- Sys.getenv("OMPI_MCA_ns_nds_vpid")
if (ndsvpid == "0") { # are we master?
  cl <- makeMPIcluster()
} else { # or are we a slave?
  sink(file="/dev/null")
  slaveLoop(makeMPImaster())
  q()
}
clusterEvalQ(cl, library(RDieHarder))
res <- parLapply(cl, c("mt19937", "mt19937_1999",
  "mt19937_1998", "R_mersenne_twister"),
  function(x) {
    dieharder(rng=x, test="operm5",
              psamples=100, seed=12345,
              rngdraws=100000)
  })
stopCluster(cl)
```

A complete example cont.

This uses `RDieHarder` to test four Mersenne-Twister implementations at once.

A simple analysis shows the four charts and prints the four p -values:

```
pdf("/tmp/snowRDH.pdf")
lapply(res, function(x) plot(x))
dev.off()

print( do.call(rbind,
               lapply(res, function(x) { x[[1]] } )))
```

A complete example cont.

```
$ salloc -w ron,joe -n 5 snowRDieharder.r
salloc: Granted job allocation 10
      [,1]
[1,] 0.1443805247
[2,] 0.0022301018
[3,] 0.0001014794
[4,] 0.0061524281
sall: Relinquishing job allocation 10
```

Example summary

We have seen

- ▶ how `littler` can help us script R tasks
- ▶ how `Rmpi`, `snow` and `slurm` can interact nicely
- ▶ a complete example using `RDieHarder` to illustrate these concepts

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Extending physical RAM limits

Two fairly recent CRAN packages ease the analysis of *large* datasets.

- ▶ `ff` which maps R objects to files and is therefore only bound by the available filesystem space
- ▶ `bigmemory` which maps R objects to dynamic objects not managed by R

Both packages can use the `biglm` package for out-of-memory (generalized) linear models.

Also worth mentioning are the older packages `g.data` for delayed data assignment from disk, `filehash` which takes a slightly more database-alike view by 'attaching' objects that are still saved on disk, and `R.huge` which also uses the disk to store the data.

biglm

The `biglm` package provides a way to operate on 'larger-than-memory' datasets by operating on 'chunks' of data at a time.

```
make.data <- function ... # see 'help(bigglm)
dataurl <-
  "http://faculty.washington.edu/tlumley/NO2.dat"
airpoll <- make.data(dataurl, chunksize=150, \
  col.names=c("logno2", "logcars", "temp", \
  "windsp", "tempgrad", "winddir", "hour", "day"))
b <- bigglm(exp(logno2)~logcars+temp+windsp, \
  data=airpoll, family=Gamma(log), \
  start=c(2,0,0,0),maxit=10)

summary(b)
```

ff

`ff` was the winner of the UseR! 2007 'large datasets' competition. It has undergone a complete rewrite for version 2.0 which was released in 2008. The following illustration (from an example in the package) is for version 1.0 of `ff` and will not run with version 2.*.

```
data("trees")
# create ffm object, convert 'trees', creates files
m <- ffm("foom.ff", c(31, 3))
m[1:31, 1:3] <- trees[1:31, 1:3]
# create a ffm.data.frame wrapper around ffm object
ffmdf <- ffm.data.frame(m, c("Girth", "Height", "Volume"))
# define formula and fit the model
fg <- log(Volume) ~ log(Girth) + log(Height)
ffmdf.out <- bigglm(fg, data=ffmdf,
                   chunksize=10, sandwich=TRUE)
```

Running `object.size()` on the `ff` object shows that it occupies less memory than the (puny) `trees` dataset.

bigmemory

The `bigmemory` package allows us to allocate and access memory managed by the operating system but 'outside' of the view of R.

```
> object.size( big.matrix(1000,1000, "double") )
```

```
[1] 372
```

```
> object.size( matrix(double(1000*1000), ncol=1000) )
```

```
[1] 8000112
```

Here we see that to R, a `big.matrix` of 1000×1000 elements occupies only 372 bytes of memory. The actual size of 800 mb is allocated by the operating system, and R interfaces it via an 'external pointer' object.

bigmemory cont.

We can illustrate `bigmemory` by adapting the previous example:

```
bm <- as.big.matrix(as.matrix(trees), type="double")
colnames(bm) <- colnames(trees)
fg <- log(Volume) ~ log(Girth) + log(Height)
bm.out <- biglm.big.matrix(fg, data=bm, chunksize=10, \
                           sandwich=TRUE)
```

As before, the memory use of the new 'out-of-memory' object is smaller than the actual dataset as the 'real' storage is outside of what the **R** memory manager sees.

`bigmemory` can also provide shared memory allocation: one (large) object can be accessed by several **R** processes as proper locking mechanisms are provided.

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

littler

Both `r` (from the `littler` package) and `Rscript` (included with `R`) allow us to write simple scripts for repeated tasks.

```
#!/usr/bin/env r
# a simple example to install one or more packages
if (is.null(argv) | length(argv)<1) {
  cat("Usage: installr.r pkg1 [pkg2 pkg3 ...]\n")
  q()
}
## adjust as necessary, see help('download.packages')
repos <- "http://cran.us.r-project.org"
lib.loc <- "/usr/local/lib/R/site-library"
install.packages(argv, lib.loc,
                  repos, dependencies=TRUE)
```

If saved as `install.r`, we can call it via

```
$ install.r ff bigmemory
```

The `getopt` package makes it a lot easier for `r` to support command-line options.

Rscript

`Rscript` can be used in a similar fashion.

Previously we had to use

```
$ R --slave < cmdfile.R  
$ cat cmdfile.R | R --slave  
$ R CMD BATCH cmdfile.R
```

or some shell-script variations around this theme.

By providing `r` and `Rscript`, we can now write 'R scripts' that are executable. This allows for automation in cron jobs, Makefile, job queues, ...

RPy

RPy packages provides access to R from Python:

```
1 from rpy import *
2 set_default_mode(NO_CONVERSION) # avoid automatic conversion
3 r.library("nnet")
4 model = r("Fxy~x+y")
5 df = r.data_frame(x = r.c(0,2,5,10,15)
6                   ,y = r.c(0,2,5,8,10)
7                   ,Fxy = r.c(0,2,5,8,10))
8 NNModel = r.nnet(model, data = df
9                  , size =10, decay =1e-3
10                 , lineout=True, skip=True
11                 , maxit=1000, Hess =True)
12 XG = r.expand_grid(x = r.seq(0,7,1), y = r.seq(0,7,1))
13 x = r.seq(0,7,1)
14 y = r.seq(0,7,1)
15
16 set_default_mode(BASIC_CONVERSION) # automatic conv. back on
17 fit = r.predict(NNModel,XG)
18 print fit
```

Outline

Motivation

Preliminaries

Measuring and profiling

RProf

RProfmem

Profiling

Faster: Vectorisation and Compiled Code

Vectorisation

Ra

Blas

GPUs

Compiled Code Overview

Inline

Rcpp

Debugging

Parallel execution: Explicitly and Implicitly

Explicitly parallel using clustered computing

Resource management and queue system

Implicitly parallel using several cores

Example

Out-of-memory processing

overview

biglm

ff

bigmemory

Automation and scripting

littler

RPy

Summary

Wrapping up

In this tutorial session, we have

- ▶ seen several ways to profile execution times;
- ▶ looked at different vectorisation examples, as well as speed increases from using compiled code;
- ▶ provided an introduction to parallel execution frameworks such as `NWS`, `MPI` and `snow` as well as the `slurm` resource managers;
- ▶ briefly looked at packages such as `ff` and `bigmemory` that can help with larger data sets;
- ▶ briefly looked at ways to script R tasks using `littler` and `Rscript`

Wrapping up

More questions ?

- ▶ A good new resource is the mailing list `r-sig-hpc`.
- ▶ And don't hesitate to email me at `dirk@eddelbuettel.com`