

# State-of-the-art in Parallel Computing with R

Markus Schmidberger  
([schmidb@ibe.med.uni-muenchen.de](mailto:schmidb@ibe.med.uni-muenchen.de))

The R User Conference 2009  
July 8-10, Agrocampus-Ouest, Rennes, France



# The Future is Parallel


Prof. Bill Dally, Nvidia, 01-2009

Thilo Kielmann, University of Amsterdam, 12-2008

# ITRS Roadmap 2005 and 2007

International Technology Roadmap for Semiconductors

VU university amsterdam



[D. Patterson, USENIX 2008 keynote]

# New Paper

submitted in December – State-of-the-art at the end of 2008



## *Journal of Statistical Software*

MMMMMM YYYY, Volume VV, Issue II.

<http://www.jstatsoft.org/>

- State of development
- Technology
- Fault-Tolerance & Load balancing
- Usability
- Acceptance
- Performance

### State-of-the-art in Parallel Computing with R

Markus Schmidberger

Ludwig-Maximilians-Universität  
München, Germany

Martin Morgan

Fred Hutchinson Cancer  
Research Center, WA, USA

Dirk Eddelbuettel

Debian Project,  
Chicago, IL, USA

Hao Yu

University of Western Ontario,  
ON, Canada

Luke Tierney

University of Iowa,  
IA, USA

Ulrich Mansmann


Ludwig-Maximilians-Universität  
München, Germany

Preprint: <http://epub.ub.uni-muenchen.de/8991/>

# Parallel Program Design

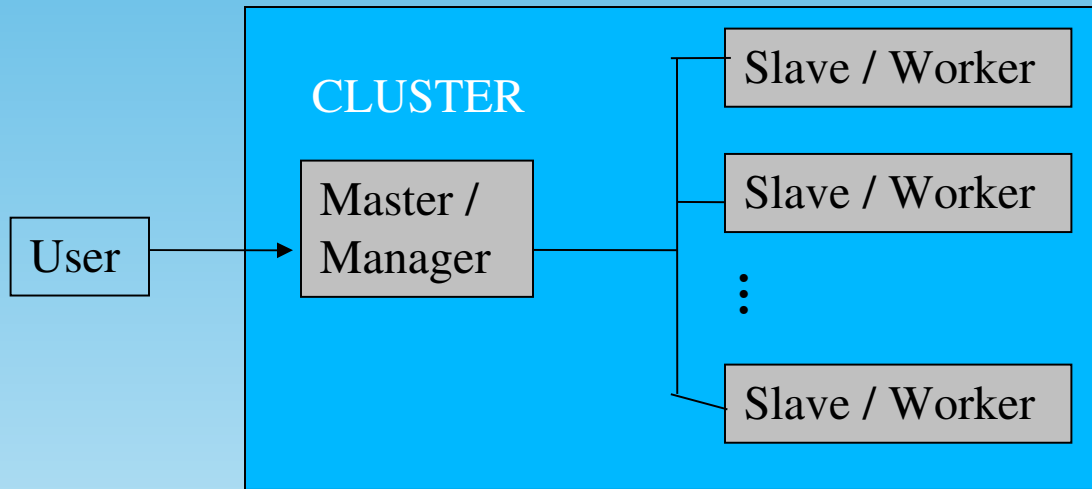
- Convert serial programs into parallel programs
  - compiler or pre-processor
  - prefer manual to automatic parallelization
    - wrong results may be produced,
    - performance may actually degrade,
    - much less flexible than manual parallelization,
    - code is too complex for automatic parallelization, etc..
- Very manual process of identifying and implementing parallelism
- Analysing the serial Code
  - understand serial code
  - profilers and performance analysis tools exist
  - identify program's hotspots or bottlenecks.
  - In the R language:
    - profile R code for memory use and evaluation time
    - ?Rprof

# Parallelization



- Multiprocessors
  - the use of two or more central processing units (CPUs) within a single computer system
  - Today: Two/Four-processors are becoming a standard for workstations
- Multicomputers
  - different parts of a program run simultaneously on two or more computers that are communicating with each other over a network
  - Computer, network, software
  - Cluster, Grid

# Master-Slave Architecture



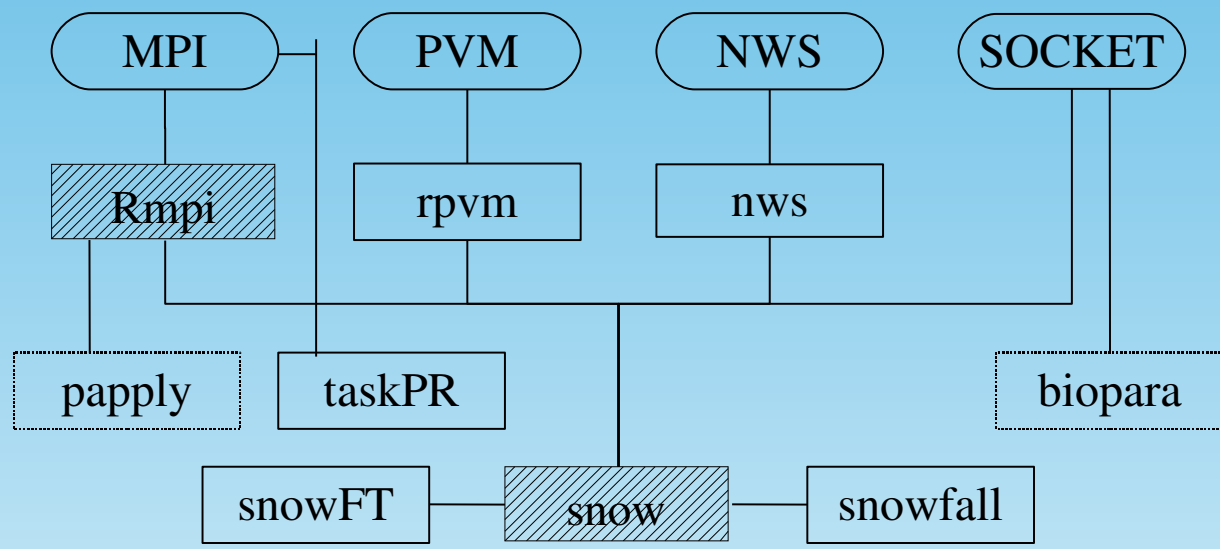
- Works on computer clusters, on multiprocessor machines and in grid computing
- You need underlying technology for communication
  - MPI: Message Passing Interface
  - PVM: Parallel Virtual Machine
  - socket, ssh
  - ( NWS: NetWorkSpace )

# Parallel R Packages

Package	Version	Websites	Technology
<i>Computer Cluster</i>			
Rmpi	0.5-6	<a href="http://cran.r-project.org/web/packages/Rmpi">http://cran.r-project.org/web/packages/Rmpi</a> <a href="http://www.stats.uwo.ca/faculty/you/Rmpi">http://www.stats.uwo.ca/faculty/you/Rmpi</a>	MPI
rpvm	1.0.2	<a href="http://cran.r-project.org/web/packages/rpvm">http://cran.r-project.org/web/packages/rpvm</a> <a href="http://www.biostat.umn.edu/~nali/SoftwareListing.html">http://www.biostat.umn.edu/~nali/SoftwareListing.html</a>	PVM
nws	1.7.0.0	<a href="http://cran.r-project.org/web/packages/nws">http://cran.r-project.org/web/packages/nws</a> <a href="http://nws-r.sourceforge.net">http://nws-r.sourceforge.net</a>	NWS and socket
snow	0.3-3	<a href="http://cran.r-project.org/web/packages/snow">http://cran.r-project.org/web/packages/snow</a> <a href="http://www.cs.uiowa.edu/~luke/R/cluster">http://www.cs.uiowa.edu/~luke/R/cluster</a>	Rmpi, rpvm, nws, socket
snowFT	0.0-2	<a href="http://cran.r-project.org/web/packages/snowFT">http://cran.r-project.org/web/packages/snowFT</a>	rpvm, snow
snowfall	1.60	<a href="http://cran.r-project.org/web/packages/snowfall">http://cran.r-project.org/web/packages/snowfall</a> <a href="http://www.imbi.uni-freiburg.de/parallel">http://www.imbi.uni-freiburg.de/parallel</a>	snow
papply	0.1	<a href="http://cran.r-project.org/web/packages/papply">http://cran.r-project.org/web/packages/papply</a> <a href="http://math.acadiau.ca/ACMMaC/software/papply.html">http://math.acadiau.ca/ACMMaC/software/papply.html</a>	Rmpi
biopara	1.5	<a href="http://cran.r-project.org/web/packages/biopara">http://cran.r-project.org/web/packages/biopara</a> <a href="http://hedwig.mgh.harvard.edu/biostatistics/node/20">http://hedwig.mgh.harvard.edu/biostatistics/node/20</a>	socket
taskPR	0.31	<a href="http://cran.r-project.org/web/packages/taskPR">http://cran.r-project.org/web/packages/taskPR</a> <a href="http://users.ece.gatech.edu/~gte810u/Parallel-R">http://users.ece.gatech.edu/~gte810u/Parallel-R</a>	MPI (only LAM/MPI)
<i>Grid Computing</i>			
GridR	0.8.4	<a href="http://cran.r-project.org/web/packages/GridR">http://cran.r-project.org/web/packages/GridR</a> <a href="http://www.stefan-rueping.de">http://www.stefan-rueping.de</a>	web service, ssh, Condor, Globus
multiR	-	<a href="http://e-science.lancs.ac.uk/multiR">http://e-science.lancs.ac.uk/multiR</a>	3 tier client/server architecture
Biocep-R	NA	<a href="http://biocep-distrib.r-forge.r-project.org">http://biocep-distrib.r-forge.r-project.org</a>	java 5
<i>Multi-core System</i>			
pnmath(0)	0.2	<a href="http://www.cs.uiowa.edu/~luke/R/experimental">http://www.cs.uiowa.edu/~luke/R/experimental</a>	openMP, Pthreads
fork	1.2.1	<a href="http://cran.r-project.org/web/packages/fork">http://cran.r-project.org/web/packages/fork</a>	Unix: fork
R/Parallel	0.6-20	<a href="http://www.rparallel.org">http://www.rparallel.org</a>	C++, file
romp	0.1a	<a href="http://code.google.com/p/romp">http://code.google.com/p/romp</a>	openMP
multicore	0.1-3	<a href="http://cran.r-project.org/web/packages/multicore">http://cran.r-project.org/web/packages/multicore</a>	fork



# Computer Cluster R Packages



- XXX R Package
- XXX Technology
- XXX No longer maintained

# Performance evaluation of R packages for computer clusters

	Component 1	Component 2	Component 3
Rmpi	0.91	1.83	3.00
nws	14.87	3.49	3.03
snow			
MPI	10.33	2.00	2.96
PVM	4.13	1.01	2.93
NWS	8.68	1.58	2.98
Socket	3.47	0.93	2.90
snowfall			
MPI	13.99	2.07	2.93
PVM	4.29	0.94	2.93
NWS	8.82	1.62	2.99
Socket	3.77	1.00	2.91

Component 1: Sending Data from the Master to all Slaves (matrix 500 x 500)

Component 2: Distributing a List of Data from the Master to the Slaves (list of matrices 500 x 500)

Component 3: Compute integral of a three-dimensional function (10.000 points)

# Performance evaluation of R packages for computer clusters


	Component 1	Component 2	Component 3
Rmpi	29.1	18.6	21.9
nws	97.3	34.8	21.2
snow			
MPI	103.2	20.1	20.5
PVM	41.2	10.1	20.5
NWS	86.7	16.0	20.8
Socket	34.8	9.3	20.2
snowfall			
MPI	109.6	20.9	20.5
PVM	43.0	9.9	20.6
NWS	88.0	16.3	20.9
Socket	37.1	9.9	20.3

# Performance - Sudoku

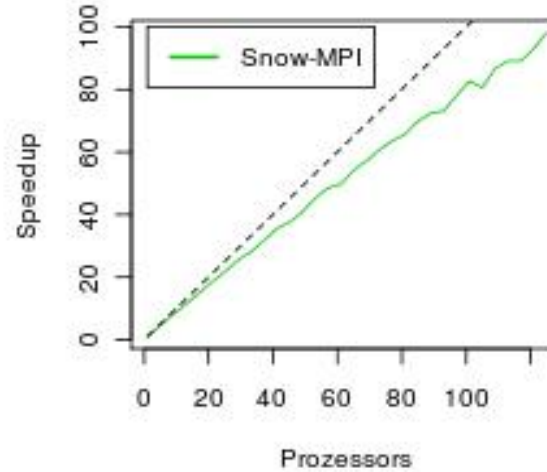
- R package: **sudoku\_2.2**
- Generates, plays, and solves Sudoku puzzles.
- Solve 10.000 Sudokus
- Distribute Sudokus equally to all nodes
- The basic rules of Sudoku are used to fill in missings, then elimination is used to find the TRUE's. If that approach runs out of steam, a guess is made and the program recurses to find either a solution or an inconsistency.

# Performance - Sudoku

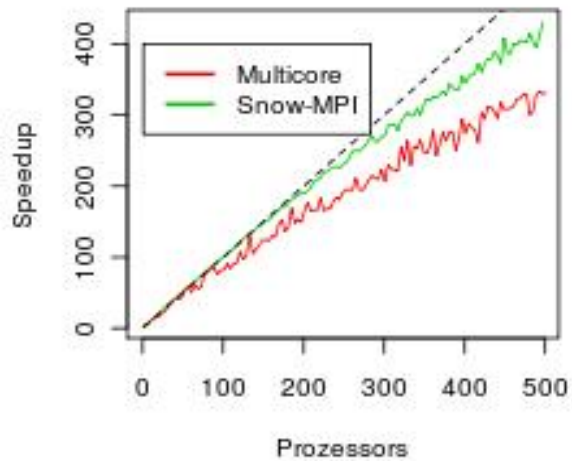
## Multicore-Machine



## IBE Cluster



## HLRB2



# State of the Art in Parallel Computing with R

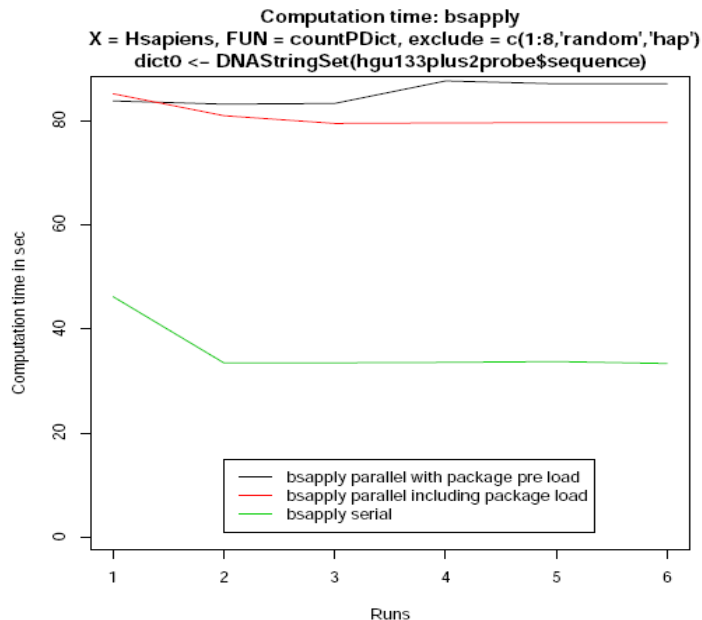
- Computer Cluster: **Rmpi** and **snow**
  - acceptable usability, wide spectrum of functionality, good performance
  - Other packages: Usability  $\leftrightarrow$  lower functionality
- Multi-core: in development
  - **Multicore** package  $\leftrightarrow$  Windows ?
  - external and architecture optimized libraries (PBLAS)
    - bottleneck in statistical computation?
  - Multicomputer packages: **Rmpi** and **snow**
    - every R instance requires its own main memory!
- Grid Computing: early-stage packages

# Which package should I use?

- Depends on your available hardware:
- Multicore machine: **multicore**
- Cluster environment:
  - \_ **Snow(-fall)** with the available communication mechanism (MPI mostly used)
  - \_ **NWS**, if you have a lot of global variables
  - \_ **Rmpi**, for excellent programmer and for high end optimization
- Grid Computing: **gridR**, which statistical application is usefull for grid computing?

# Tips for Parallel Computing

- Communication is much slower than computation.
  - functions produce large results, reduce results on the worker before returning.
  - additional function parameters can be huge.



## bsapply and countPDict Example

```
R> params <- new("BSPParams", X = Hsapiens,  
  FUN = countPDict)  
R> library(hgu133plus2probe)  
R> dict0 <- DNASTringSet(hgu133plus2probe$sequence)  
R> pdict0 <- PDict(dict0)  
R> bsapply(params, pdict = pdict0)
```



# Tips for Parallel Computing

- Random Generators have to be used with care; special-purpose packages **rsprng**, **rlecuyer** (and **snow**) are available.

```
R> clusterCall(cl, runif, 3)

[[1]]
[1] 0.4351672 0.7394578 0.2008757

[[2]]
[1] 0.4351672 0.7394578 0.2008757

...

[[10]]
[1] 0.4351672 0.7394578 0.2008757
```

```
R> clusterSetupSPRNG(cl)
R> clusterCall(cl, runif, 3)

[[1]]
[1] 0.014266542 0.749391854 0.007316102

[[2]]
[1] 0.8390032 0.8424790 0.8896625

...

[[10]]
[1] 0.591217470 0.121211511 0.002844222
```

- lexical scoping: requires some care to avoid transmitting unnecessary data to workers
  - Functions used in apply-like calls should be defined in the global environment, or in a package name space.

# HELP

- „State-of-the-Art in Parallel Computing with R“; Schmidberger, et.al.; JSS 2009
- CRAN Task View 'High Performance and Parallel Computing'
  - <http://cran.r-project.org/web/views/HighPerformanceComputing.html>
- R Mailinglist 'R SIG on High-Performance Computing'
  - <https://stat.ethz.ch/mailman/listinfo/r-sig-hpc>

# Conclusion & Future

- Parallel Computing can help improving performance,
  - but first of all improve your serial code (profiling, vectorization, ...).
  - but be careful with communication costs.
- First Parallel Implementations are easy,
  - but there are a lot of stumbling blocks.
- Parallel Computing with R needs to be improved:
  - Teach R users to think in parallel
  - Integration of R code into multi-core environments
  - Cloud Computing with R
  - Computing power of graphic processing units
- Flexibility of R package system allows integration of many different technologies.

# Acknowledgment

## Parallel R

Martin Morgan

Dirk Eddelbuettel

Hao Yu

Luke Tierney

Anthony Rossini

## LRZ

HPC Team

Ferdinand Jamitzky

200.000 CPUh

## AffyPara Package

Ulrich Mansmann

Esmeralda Vicedo

Klaus Rüschoer

Robert Gentleman's Group

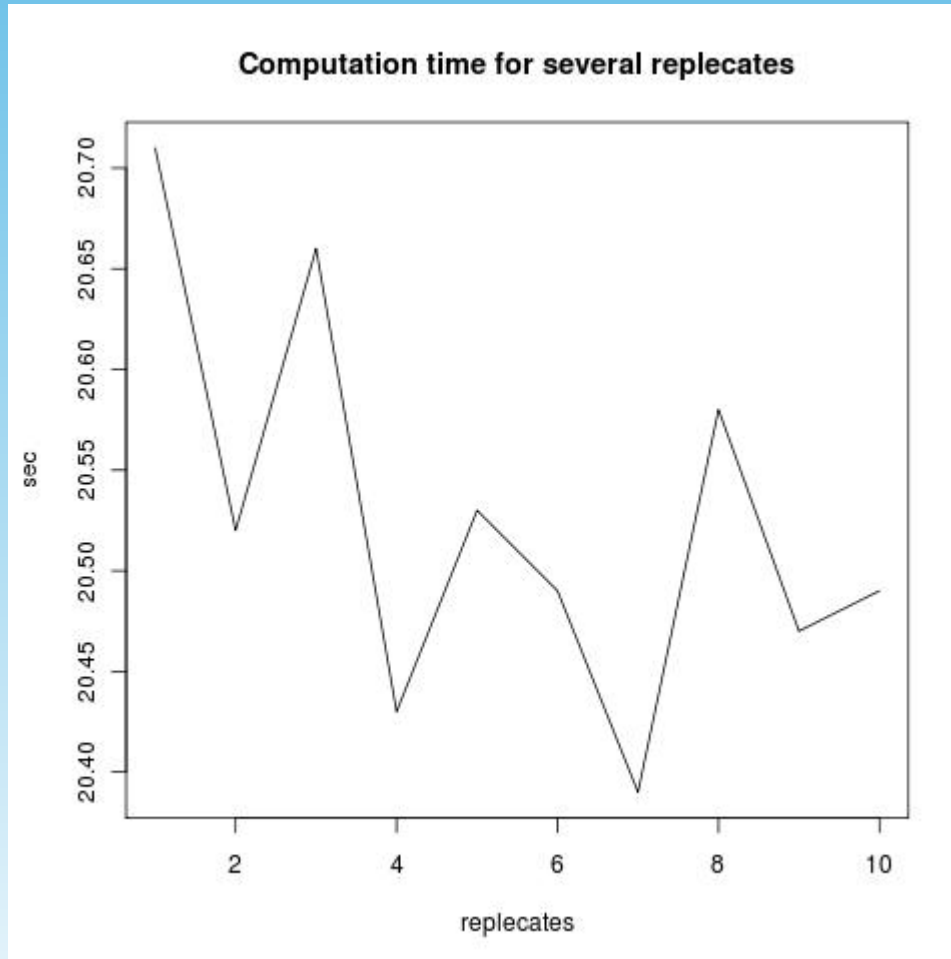
Dipl.-Tech. Math. Markus Schmidberger  
schmidb@ibe.med.uni-muenchen.de  
<http://ibe.med.uni-muenchen.de>

**Thanks for your  
attention**

# State of the Art in Parallel Computing with R

	Learnability	Efficiency	Memorability	Errors	Satisfaction
Rmpi	+	--	++	+	+
rpvm	--	--	+	-	--
nws	+	+	++	+	+
snow	+	++	++	+	++
snowFT	+	++	++	+	++
snowfall	++	++	++	+	++
papply	+	+	++	+	0
biopara	--	--	0	0	--
taskPR	+	-	++	0	-
	Learnability	Efficiency	Memorability	Errors	Satisfaction
<i>Grid Computing</i>					
gridR	+	-	+	-	-
R/parallel	0	+	+	-	0
<i>Multi-core System</i>					
pnmath(0)	++	++	++	++	+
fork	+	-	+	-	-
<b>Multicore</b>	<b>+</b>	<b>++</b>	<b>++</b>	<b>0</b>	<b>+</b>

# Computation Time for 10 replicates



# Simple Parallelization

```
L <- list(a=c(1:10), b=c(2:12), c=c(4:14))
```

## Seriell:

```
for(i in 1:3) { res[[i]] <- mean( L[[i]] ) }
```

```
res <- lapply(L, mean)
```



## Parallel:

```
library(snow)
```

```
cl <- makeCluster(3, type='SOCK')
```

```
res <- clusterApply(cl, L, mean)
```

```
stopCluster(cl)
```



# Simple Parallelization for Statisticians

- Bootstrapping: time-consuming and simple to parallelize
- library(**boot**): generating bootstrap replicates
- Example: generalized linear model fit for data on the cost of constructing nuclear power plants. 999 bootstraps
- Serial: 9.2 sec <-> 3 nodes: 3.1 sec