

Running ISA in parallel with the **snow** package

Gábor Csárdi

September 20, 2013

1 Running ISA in parallel

In this document we show a little example on how to speed up ISA analysis [?, ?, ?] by running the ISA iterations in parallel, on a computer cluster, or a multiprocessor machine.

Since a typical ISA analysis consists of using a range of row/column thresholds and these runs are independent of each other; it is trivial to parallelize the task by performing the iterations for different threshold parameters on different processors or computers. Here we show an example on how to do this easily with the **snow** [?] and the **Rmpi** [?] packages:

```
> library(isa2)
> library(snow)
> library(Rmpi)
```

We generate some simple in-silico data.

```
> pdata <- isa.in.silico()
```

Next, we create the MPI cluster, with eight working nodes. You need to have a working MPI installation for this. See more in the documentation of the **Rmpi** and **snow** packages.

```
> clu <- makeMPIcluster(8)
```

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

Now we load the **isa2** package on all nodes and distribute the input data to them.

```
> invisible(clusterEvalQ(clu, library(isa2)))
> clusterExport(clu, "pdata")
```

Next, we create a big matrix in which each row is a combination of the threshold parameters. This will be needed for the parallel run.

```
> thr <- seq(1,3,by=0.2)
> thr.list <- expand.grid(thr, thr)
```

First we run the ISA on a single processor only, and measure the running time.

```
> system.time(modules <- isa(pdata[[1]], thr.row=thr, thr.col=thr))

   user  system elapsed 
45.029   3.368  50.852
```

Let us now do a parallel run, again, with measuring the running time. If you are really running this on multiple CPUs, then it is much faster.

```
> system.time(modules.par <- parApply(clu, thr.list, 1, function(x) {
  isa(pdata[[1]], thr.row=x[1], thr.col=x[2])
}))

   user  system elapsed 
0.571   3.493  19.579
```

Finally, stop the cluster.

```
> stopCluster(clu)

[1] 1
```

2 Session information

The version number of R and packages loaded for generating this vignette were:

- R version 3.0.1 (2013-05-16), x86_64-apple-darwin10.8.0
- Locale:
en_US.UTF-8/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
- Base packages: base, datasets, graphics, grDevices, methods, stats, utils
- Other packages: isa2 0.3.3, Rmpi 0.6-3, snow 0.3-12
- Loaded via a namespace (and not attached): tools 3.0.1

References

- [Bergmann et al., 2003] Bergmann, S., Ihmels, J., and Barkai, N. (2003). Iterative signature algorithm for the analysis of large-scale gene expression data. *Phys Rev E Nonlin Soft Matter Phys*, page 031902.
- [Csárdi, 2009] Csárdi, G. (Apr 1, 2009). *isa2: The Iterative Signature Algorithm*. R package version 0.1.

- [Ihmels, 2002] Ihmels, J., Friedlander, G., Bergmann, S., Sarig, O., Ziv, Y., Barkai, N. (2002). Revealing modular organization in the yeast transcriptional network. *Nat Genet*, page 370–7.
- [Ihmels, 2004] Ihmels, J., Bergmann, S., Barkai, N. (2004). Defining transcription modules using large-scale gene expression data. *Bioinformatics*, page 1993–2003.
- [Tierney, 2009] Luke Tierney, A. J. Rossini, Na Li and H. Sevcikova (2009). *snow: Simple Network of Workstations*. R package version 0.3-3.
- [Yu, 2007] Hao Yu (2007). *Rmpi: Interface (Wrapper) to MPI (Message-Passing Interface)*. R package version 0.5-5.