

**SOMS Friday Seminar
April 18, 2008
SMC 305
2:30-3:25**

**Stalking the Interactive Terabyte with R: Data-Parallel
Statistical Computing**

**George Ostrouchov
Senior Research Staff Member
Oak Ridge National Laboratory**

ABSTRACT

Moore's law of transistor density no longer translates into increasing processor clock rates and instead we see increases in the number of processor cores. It is no longer effective to "just wait" for the next faster machine to see a favorite algorithm process more data faster and the need for data-parallel algorithms is more pressing. I will describe what this means for statistical computing and what are the current efforts to address this challenge. Data parallelism is a necessary component in answering this challenge but it often requires rethinking standard statistical algorithms. I will include some experiences on medium size cluster computers and on supercomputers while developing a runtime environment for R and its Rmpi package.

Refreshments provided after seminar