

Experience from AMS MonteCarlo Mass–Production

**A.Biland / ETHZ
CERN, 18.10.02**

**usually, a production–cluster
consists of several/many
computing nodes and
few fileserver**

==>

fileserver(s) are bottleneck

avoid writing via NFS/AFS

**write ntuples to local
scratch–disk on comp.
node, gzip them and
transfer to fileserver
via rcp/rsync/...**


**during 2001/2002 mass
production, more than
1TB of data /month
produced at ETH
(~40x 1GHz machines)**

**where to store files ???
several 100GBs lost in
disk-crashes etc. at CERN
(re-re-reruns needed)**

**file-transfer ETH-CERN:
bottleneck seems to be
inside CERN !!!**


**Key–Problem in all kinds
of parallel MC production:
rndm–synchronization**

**'normal experiments':
less severe because each
event needs many rndms**



**2 streams start using same
rndm num., but no synchr.**

**problem in AMS:
long sequence of 0–events**



==> full event synchr. !!!

possible solutions:

**use real random numbers
instead of pseudo-rndm !**

- not always available**
- impossible to rerun**

**use slightly different
p-distributions for
different streams**

**==> will de-synchronize
after some events**

- ok for fixed-p, but
for 'real-p dist.' ?**
- for >100 streams ??**

possible solutions:

**use seeds well enough
seperated**

**how many rndm-numbers
does a stream need ????**

**/afs/cern.ch/users/b/biland/
public/AMS/rndm.seeds**

**contains (at the moment)
>200'000 seeds, all
separated by 1'000'000'000
rndm numbers**

a wish to Vitaly:

**would be helpful if each
job would write a short
file (that can be easily
parsed/sourced by scripts)
containing key–infos
to continue the job
(next evt.nr., rndm–seed,...)
and some job statistics
(#ev. proc, #ev. written,...)**

**none–existence of this
file would also indicate
that job crashed...**