

KLOE software

Old KLOE computing relied on outdated IBM infrastructure,
was impossible to run simulation/reconstruction outside of the KLOE IBM farm

Significant effort was invested to migrate the software code to the modern Linux environment.

New release of reconstruction/simulation version adopted for Linux is ready to use:

- × Based on AIX original last KLOE-1 reconstruction/simulation codes (**no KLOE-2 or Kpm part yet**)
- × Many issues was solved during migration: to Intel x86 from Power8 CPU (little/big endian), x86_64/x32 addressing, gfortran from IBM xlf compiler, AIX OS functions (threading, signals, etc), makefiles
- × Numerous bugs were fixed: memory access related, NaNs, infinite loops, etc...
- × Improved stability of tracking, propagation to beam point, vertexing
- × Simulation was extended with new generator plugins, MCtruth information, etc
- × Extensive tests for Linux vs AIX versions consistency

Hosted at the git repositories (mirrored):

<https://gitlab.ph.liv.ac.uk/kloe-hvp/soft/>

<https://baltig.infn.it/kloe>

with all CVS history preserved from March 1996


DB2 database is hosted now in VM
(with calibrations, runs, etc info)

Anyone from the collaboration can install and run KSoft following the instruction
(including DB2 database deployment)


Linux vs AIX

Linux DBV-50-rc1
vs AIX modified "production" build
(with some backported modifications)


Tracks




consistency $\sim 10^{-6}$




Vertices




consistency $\sim 10^{-5}$



Clusters



consistency $\sim 10^{-7}$



Fluctuations are from single-point floating precision used in the original code

Processing time/ disk usage

Available CPU resources:

IBM KLOE cluster - 304 Power8 CPU (~3 times slower compared to modern Intel)

will be used only for cross checks

Liverpool cluster - 288 CPU (but shared between all in Physics Department, usually in idle)

CNAF - 60 CPU dedicated for KLOE

Full KLOE-1 statistics reprocessing time:

Liverpool ~ 3 weeks

on IBM ~ 2. months

in CNAF ~ 3.5 months

+ ~ same order for MC production
(new generators, ...)

Overall disk space requirement:

150 TB KLOE-1 raw data

+ **300 TB ntuples for $\pi\pi\gamma$ group:**

(as minimum for one version of reprocessing)

~100 TB full prod2root ntuples

~few 10 TB shrunk fast ntuples

~100 TB MC production

... etc

If to use only allocated CNAF resources (60 CPU) than data + MC **will take ~ 1 year....**
will be very problematic for us

Is there a possibility to increase CPU at CNAF?

Modern approach also involves the use of distributed resources across different computing centers

Ready to start re-processing in January 2026

(completeness of data crucially rely on recovery from the tape library)