



# Software Usage on Cray Systems across Three Centers (NICS, ORNL and CSCS)

Bilel Hadri, Mark Fahey, Timothy Robinson,  
and William Renaud

CUG 2012, May 3<sup>rd</sup>, 2012



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre



# Contents

- **Introduction and Motivations**
- **Overview on ALTD Tool**
- **Data Mining over 3 centers ( Kraken, Jaguar, Rosa)**
  - Linkline
  - Execution
- **Conclusions and Future Work**

# Contents

- **Introduction and Motivations**
- Overview on ALTD Tool
- Data Mining over 3 centers ( Kraken, Jaguar, Rosa)
  - Linkline
  - Execution
- Conclusions and Future Work

# Software/library/applications

Kraken, Jaguar and Rosa support:

- **Several categories of software/library**

- Linear algebra
- I/O
- Performance tools
- Debugger
- Chemistry
- Molecular dynamic
- Materials
- Communications
- Visualization

- **Multiple versions**

- hdf5 (1.6.10 - 1.8.3 - 1.8.4 - 1.8.5 - 1.8.6)
- netcdf (3.6.2 - 4.0 - 4.0.1 - 4.1 - 4.1.1.0 4.1.2 4.1.3/)
- libsci ( 10.4.5/ 10.5.0/ 10.5.01/ 10.5.02/ 11.0.00/ 11.0.01 11.0.02/ 11.0.04 11.0.06/)

- **Multiple builds with different compiler:**

- example with FFTW:

cnl3.1\_cray7.4.3/      cnl3.1\_intel12.0.4.191/      sles11.1\_cray7.4.3/      sles11.1\_intel12.0.4.191/  
cnl3.1\_gnu4.6.1/      cnl3.1\_pgi11.6.0/      sles11.1\_gnu4.6.1/      sles11.1\_pgi11.6.0/



# Issue

- As of today:
  - Rule of the thumbs from the staff:
    - not strictly accurate and reliable
  - Modulefile, logs, Surveys:
    - Incomplete data

“Questions were raised around whether other techniques such as user survey, and gathering of library requirements a priori rather than as forensics would be more cost effective and achieve similar goals.”

**→ Solution ALTD: Automatic Library Tracking Database**

# Contents

- Introduction and Motivations
- **Overview on ALTD Tool**
- Data Mining over 3 centers ( Kraken, Jaguar, Rosa)
  - Linkline
  - Execution
- Conclusions and Future Work

# Objectives and Goals

- **A primary objective of ALTD :**
  - track only libraries linked into the applications (not the function calls)
  - track parallel executables launched (how often are the libraries used?)
  
- **Have as little impact on user as possible**
  - Lightweight solution
    - No runtime increase
    - Only link time and job launch have marginal increase in time
  - Do not change user experience
    - Linker and job launcher work as expected
  
- **Intercept the whole library path to retrieve valuable information on :**
  - Package name
  - Version number
  - Build configuration



# ALTD design

- Intercepting the GNU linker (ld) to get the linkage information
  - Intercepting the job launcher (aprun)
- **Wrapping the linker and the job launcher through scripts is a simple and efficient way to obtain the information automatically and transparently.**
- **ld - Intercept link line**
    - Update tags table
    - Call real linker (with tracemap option)
    - Use output from tracemap to find libraries linked into executable
    - Update linkline table
  - **aprun- Intercept job launcher**
    - Pull information from ALTD section header in executable
    - Update jobs table
    - Call real job launcher
- **Storing information about compilation and execution into a database that can be mined to provide reports.**

# ALTD database results

- ALTD generates records into 3 tables:
  - Tags: entry for every link executed
  - Linkline: entry for each unique link line
  - Jobs: entry for each executable launched

linkline_id	linkline
14437	./bin/cg.B.4 /usr/lib/./lib64/crti.o /usr/lib/./lib64/crti.o /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtbegin.o /sw/xt/tau/2.19/cnl2.2_gnu4.4.1/tau-2.19/craycnl/lib/libtau-mpi-gnu-mpi-pdt.a /sw/xt/tau/2.19/cnl2.2_gnu4.4.1/tau-2.19/craycnl/lib/libtau-mpi-gnu-mpi-pdt.a /usr/lib/./lib64/libpthread.a /opt/cray/mpt/4.0.1/xt/seastar/mpich2-gnu/lib/libmpich.a /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/./lib64/libmpich.a /usr/lib/./lib64/libmpich.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/libgfortranbegin.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/libgcc.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/libgcc_eh.a /usr/lib/./lib64/libc.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtend.o /usr/lib/./lib64/crti.o
14438	highmass3d.Linux.CC.ex /usr/lib64/crti.o /usr/lib64/crti.o /opt/pgi/9.0.4/linux86-64/9.0-4/lib/trace_init.o /usr/lib64/gcc/x86_64-suse-linux/4.1.2/crtbegin.o /sw/xt/hypre/2.0.0/cnl2.2_pgi9.0.1/lib/libHYPRE.a /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/./lib64/libmpich.a /usr/lib/./lib64/libmpich.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a /usr/lib64/libpthread.a /usr/lib64/libm.a /usr/local/lib/libmpich.a /opt/pgi/9.0.4/linux86-64/9.0-4/lib/libtd.a /opt/pgi/9.0.4/linux86-64/9.0-4/lib/libc.a /opt/pgi/9.0.4/linux86-64/9.0-4/lib/libpgf90.a /opt/pgi/9.0.4/linux86-64/9.0-4/lib/libpgc.a /usr/lib64/librt.a /usr/lib64/libpthread.a /usr/lib64/libm.a /usr/lib64/gcc/x86_64-suse-linux/4.1.2/libgcc_eh.a /usr/lib64/libc.a /usr/lib64/gcc/x86_64-suse-linux/4.1.2/crtend.o /usr/lib64/crti.o
14439	probeTest /usr/lib/./lib64/crti.o /usr/lib/./lib64/crti.o /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtbegin.o /opt/cray/mpt/4.0.1/xt/seastar/mpich2-gnu/lib/libmpich.a /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/./lib64/libmpich.a /usr/lib/./lib64/libmpich.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a /usr/lib/./lib64/libpthread.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/libgcc_eh.a /usr/lib/./lib64/libc.a /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtend.o /usr/lib/./lib64/crti.o

a) linkline table

tag_id	linkline_id	username	exit_code	link_date
91126	14437	user1	0	2010-04-28
91127	0	user2	-1	2010-04-28
91128	14435	user3	0	2010-04-28
91129	6835	user2	0	2010-04-28
91130	14438	user4	0	2010-04-28
91131	14439	user1	0	2010-04-28
91132	14439	user1	0	2010-04-28

b) tag\_id table

run_id	tag_id	executable	username	run_date	job_launch_id	build_machine
144091	91126	/nics/b/home/user1/NPB3.3/bin/cg.B.4	user1	2010-04-28	548346	kraken
144099	91131	/nics/b/home/user1/probeTest	user1	2010-04-28	548357	kraken
144102	91132	/nics/b/home/user1/probeTest	user1	2010-04-28	548357	kraken
144179	91128	/lustre/scratch/user3/CH4/vasp_vtst.x	user3	2010-04-28	548444	kraken
144192	91128	/lustre/scratch/user3/CH4/vasp_vtst.x	user3	2010-04-28	548488	kraken
144356	91128	/lustre/scratch/user5/snc/CH4/vasp_vtst.x	user5	2010-04-29	548638	kraken

c) job\_id table

# Linktable

linkline_id	linkline
14437	<b>../bin/cg.B.4</b> /usr/lib/./lib64/crt1.o /usr/lib/./lib64/crti.o /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtbeginT.o <b>/sw/xt/tau/2.19/cnl2.2_gnu4.4.1/tau-2.19/craycnl/lib/libTauMpi-gnu-mpi-pdt.a</b> <b>/sw/xt/tau/2.19/cnl2.2_gnu4.4.1/tau-2.19/craycnl/lib/libtau-gnu-mpi-pdt.a</b> /usr/lib/./lib64/libpthread.a /opt/cray/mpt/4.0.1/xt/seastar/mpich2-gnu/lib/libmpich.a /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/alps/libalpslli.a /usr/lib/alps/libalpsutil.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a [... gcc 4.4.2 libraries ...] /usr/lib/./lib64/libc.a /usr/lib/./lib64/crtn.o
14438	<b>highmass3d.Linux.CC.ex</b> /usr/lib64/crt1.o /usr/lib64/crti.o /opt/pgi/9.0.4/linux86-64/9.0-4/lib/trace_init.o /usr/lib64/gcc/x86_64-suse-linux/4.1.2/crtbeginT.o <b>/sw/xt/hypre/2.0.0/cnl2.2_pgi9.0.1/lib/libHYPRE.a</b> /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/alps/libalpslli.a /usr/lib/alps/libalpsutil.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a /usr/lib64/libpthread.a /usr/lib64/libm.a /usr/local/lib/libmpich.a [... pgi 9.0.4 libraries ...] /usr/lib64/librt.a /usr/lib64/libpthread.a /usr/lib64/libm.a /usr/lib64/gcc/x86_64-suse-linux/4.1.2/libgcc_eh.a /usr/lib64/libc.a /usr/lib64/gcc/x86_64-suse-linux/4.1.2/crtend.o /usr/lib64/crtn.o
14439	<b>probeTest</b> /usr/lib/./lib64/crt1.o /usr/lib/./lib64/crti.o /opt/gcc/4.4.2/snos/lib/gcc/x86_64-suse-linux/4.4.2/crtbeginT.o /opt/cray/mpt/4.0.1/xt/seastar/mpich2-gnu/lib/libmpich.a /opt/cray/pmi/1.0-1.0000.7628.10.2.ss/lib64/libpmi.a /usr/lib/alps/libalpslli.a /usr/lib/alps/libalpsutil.a /opt/xt-pe/2.2.41A/lib/snos64/libportals.a /usr/lib/./lib64/libpthread.a [... gcc 4.4.2 libraries ...] /usr/lib/./lib64/libc.a /usr/lib/./lib64/crtn.o

# Job table

run_inc	tag_id	executable	username	run_date	Job_launch_id	build_machine
144091	91126	/nics/b/home/user1/NPB3.3/bin/cg.B.4	user1	2010-04-28	548346	kraken
144099	91131	/nics/b/home/user1/probeTest	user1	2010-04-28	548357	kraken
144102	91132	/nics/b/home/user1/probeTest	user1	2010-04-28	548357	kraken
144179	91128	/lustre/scratch/user3/CH4/vasp_vtst.x	user3	2010-04-28	548444	kraken
144192	91128	/lustre/scratch/user3/CH4/vasp_vtst.x	user3	2010-04-28	548488	kraken
144356	91128	/lustre/scratch/user5/CH4/vasp_vtst.x	user5	2010-04-28	548638	kraken

# Contents

- Introduction and Motivations
- Overview on ALTD Tool
- **Data Mining over 3 centers ( Kraken, Jaguar, Rosa)**
  - Linkline
  - Execution
- Conclusions and Future Work

# OMG ! What are they doing !?

- Early 2010, on Kraken, a new recent tree has been put in place, `/sw/xt/` instead of `/sw/xt5/`

- End 2010, still some users are using old builds: 

```
/sw/xt5/lapack/3.1.1/cnl2.1_pgi7.2.3/lib/liblapack.a      18
/sw/xt5/zip/2.1/sles10.1_pgi7.2.3/lib/libsz.a          16
/sw/xt5/gsl/1.11/cnl2.1_pgi7.2.3/lib/libgsl.a         15
/sw/xt5/hdf5/1.8.2/cnl2.1_pgi7.2.5/lib/libhdf5.a      5
/sw/xt5/blas/ref/cnl2.1_pgi7.2.3/lib/libblas.a        4
/sw/xt5/fftw/3.1.2/sles10.1_gnu4.2.4/lib/libfftw3f.a  3
/sw/xt5/fpmpi/1.1/cnl2.1_pgi7.2.5/lib/libfpmpi_papi.a 2
/sw/xt5/hdf4/4.2r4/cnl2.1_pgi7.2.5/lib/libdf.a       1
/sw/xt5/hdf5/1.6.7/cnl2.1_pgi7.2.3/lib/libhdf5.a     1
/sw/xt5/netcdf/3.6.2/sles10.1_pgi7.2.3/lib/libnetcdf.a 1
/sw/xt5/hdf4/4.2r4/cnl2.1_pgi7.2.5/lib/libmfhdf.a    1
```

- Some users not using the modulefiles  

- Usage of reference BLAS and LAPACK on Kraken   

# ALTD : full paths

- ALTD can detect also the location of the library used and subpackages

/opt/cray/hdf5-parallel/1.8.4.1/hdf5-parallel-pgi/lib/libhdf5.a	1538	} HDF5 installed by vendor
/opt/cray/hdf5-parallel/1.8.4.1/hdf5-parallel-pgi/lib/libhdf5_fortran.a	1407	
/opt/cray/hdf5-parallel/1.8.5.0/hdf5-parallel-pgi/lib/libhdf5.a	1365	
/opt/cray/hdf5-parallel/1.8.5.0/hdf5-parallel-pgi/lib/libhdf5_fortran.a	1356	

/sw/xt/hdf5/1.6.10/cnl2.2_pgi10.4.0_par/lib/libhdf5.a	266	} HDF5 installed by staff
/sw/xt/hdf5/1.8.5/cnl2.2_pgi10.4.0/lib/libhdf5.a	230	
/sw/xt/hdf5/1.8.5/cnl2.2_gnu4.4.3/lib/libhdf5.a	192	

/nics/b/home/usr1/hdf5-1.8.7/src/.libs/libhdf5.a	37	} HDF5 installed by users
/nics/b/home/usr2/deploy/lib/libhdf5.a	32	
/nics/b/home/usr3/hdf5-1.6.5/hdf5/lib/libhdf5.a	31	
/nics/d/home/usr4hdf5patch1/lib/libhdf5_fortran.a	17	

/opt/petsc/3.1.00/real/gnu/linux/lib/44/libparmetis_quadcore.a	337	} PETSc subpackages usage
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libcraypetsc_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libpardist_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libsuperlu_dist_2.3_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libmumps_common_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libsuperlu_4.0_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libdmumps_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libmetis_quadcore.a	337	
/opt/petsc/3.1.00/real/gnu/linux/lib/44/libhypre_quadcore.a	328	Why quadcore for Istanbul ?

# Data mining

- **ALTD has recorded the following data at each site for Year 2011:**
  - **Kraken:**
    - 456,437 successful compilations by 860 users
    - 1,434,972 application executions by 919 users.
  - **Jaguar:**
    - 1,024,793 successful compilations performed by 684 users
    - 1,325,538 application executions by 671 users.
  - **Rosa:**
    - 103,451 successful compilations by 254 users,
    - 501,102 application executions by 309 users.
- **The percentage of active users who have never compiled a code is about 18% and 6%, on Rosa and Kraken, respectively.**
- **The presence of these “black-box” users, (running applications installed either by the centers’ staff or by their colleagues ) needs to be taken into account when center staff are considering the installation and maintenance of third-party applications for their users.**



# Compiler usage with MPI

## USAGE OF COMPILERS (NUMBER OF INSTANCES)

Compiler	Kraken	Jaguar	Rosa
GNU	26689	70854	<b>9407</b>
PGI	<b>51154</b>	<b>132345</b>	6116
Intel	6321	55182	1729
CCE	69	343	1415
Pathscale	14	1486	389

## USAGE OF COMPILERS (NUMBER OF USERS)

Compiler	Kraken	Jaguar	Rosa
GNU	189	190	85
PGI	<b>609</b>	<b>524</b>	<b>87</b>
Intel	146	64	41
CCE	3	3	39
Pathscale	3	38	20

# ALTD - Cray

- **Detected in 2011 that ALTD was not tracking when Cray Compiler was used:**
  - CCE uses its own linker not the GNU linker located in /usr/bin
  - /opt/cray/cce/7.3.3/cray-binutils/x86\_64-unknown-linux-gnu/bin/ld
  - Is there a reason ?
- **Fix:**
  - setenv LINKER\_X86\_64 /sw/altd/bin/ld
  - Modify the PrgEnv-Cray: unload/load ALTD
  - Ticket to CRAY : response “for user who want to use an alternative linker path, the environment variable ALT\_LINKER is created.” This environment variable is added to xt-asyncpe 5.05
- **Recent usage of Cray in 2012:**
  - Kraken : 198 out of 42649 instances with 8 unique users out of 504 total users.
  - Jaguar(TITAN): 11336 (140637) instances by 7users (373)

# Software usage during linking process

- **Top 10 most used libraries fall into three major categories (on all machines):**
  - numerical libraries (LibSci, FFTW, ACML, PETSc),
  - I/O software (HDF5, NetCDF),
  - Performance analysis tools (Craypat, PAPI, TAU).

# Numerical Libraries

Library usage ranked by number of instances and number of users on Kraken

Library	instances	users	Library/ version	instances	users
Libsci	42271	291	libsci/10.5.02	29787	220
atlas	35954	8	fftw/3.2.2.1	15987	128
fftw	24494	235	xt-libsci/10.4.5	12167	169
acml	3537	59	fftw/2.1.5	3710	64
petsc	2460	20	acml/4.4.0/	3088	39
sprng	1745	13	sprng/2.0b/	1739	12
arpack	1721	11	petsc/3.1.05	1571	13
tspl	1517	14	arpack/2008	1543	1
gsl	1451	48	tpsl/1.0.0/	1517	14
fftpack	1317	35	gsl/1.14	1063	39

 Installed by staff

Jaguar

Library	instances	users	Library/ version	instances	users
libsci	74970	317	libsci/10.4.4/	47383	245
fftw	48728	163	fftw/3.2.2.1/	44779	109
acml	17198	58	libsci/10.5.0/	26303	208
trilinos	7518	25	acml/4.3.0/	9360	43
petsc	6008	58	acml/4.4.0/	7727	32
parmetis	1810	19	petsc/3.0.0.10/	1882	29
umfpack	1773	24	fftw/3.2.2/	1832	15
arpack	1166	12	parmetis/3.1.1	1793	15
fftpack	1069	21	trilinos/10.4.0/	1786	10
pspline	1066	16	petsc/3.1.04/	1152	25

Rosa

Library	instances	users	Library/ version	instances	users
libsci	6240	109	fftw/3.2.2.1	4497	53
fftw	6042	84	libsci/11.0.01/	2806	33
acml	2020	48	libsci/10.5.02	1919	55
trilinos	1090	9	acml/4.4.0/	1123	44
tpsl	974	7	libsci/11.0.04/	974	33
parmetis	913	8	trilinos/10.6.0	731	1
umfpack	865	9	fftw/3.3.0.0	721	27
petsc	469	11	tpsl/1.0.0/	532	3
mk1	145	7	libsci/11.0.03/	512	24
Gsl	116	4	fftw/2.1.5	346	16

# Application linked with different versions of LibSci

Rosa

Kraken

One of the most powerful aspects of ALTD is its ability to identify users running codes that were built with legacy versions of libraries (particularly ones that are known to provide suboptimal performance, or are known to contain bugs)

Usage of LibSci versions				
Month	11.0.3	11.0.4	11.0.5	11.0.6
Dec 11	2326	18	0	0
Jan 12	4875	2387	0	0
Febr12	6388	4459	0	0
March 12	2280	1693	4	0
Number of users using LibSci versions				
Dec 11	33	4	0	0
Jan 12	17	29	0	0
Febr12	12	30	0	0
March 12	12	26	2	0

Usage of LibSci versions						
Month	10.4.5	10.5.0	10.5.02	11.0.01	11.0.4	11.0.6
Dec 11	2977	0	18886	1930	0	0
Jan 12	4625	4	20370	621	77	0
Febr12	2794	17	14262	692	1	0
March 12	520	0	8582	979	8890	7
Number of users using LibSci versions						
Dec 11	17	0	71	5	0	0
Jan 12	14	1	76	4	2	0
Febr12	17	1	95	7	1	0
March 12	11	0	47	7	75	2

# I/O and Performance tools

## I/O usage on Kraken

Library/ version	instances	users
iobuf/beta	5763	16
hdf5/1.8.4.1	2023	48
hdf5/1.8.5.0	1931	50
hdf5-par/1.8.5.0	1811	23
netcdf/3.6.2/	1600	25
szip/2.1	1533	51
hdf5-par/1.8.4.1	1467	17
netcdf/3.6.3/	868	5
netcdf/4.0.1.3/	593	16
hdf5/1.6.10	505	18

## Performance tools usage on Jaguar

Library/ version	instances	users
papi/3.7.2	3787	121
craypat/5.1.0/	1331	29
vampirtrace/5.11	782	16
vampirtrace/5.12	780	4
craypat/5.1.3	767	19
papi/3.6.2.2	388	7
papi/4.1.0.0.2	285	17
craypat/5.0.2/	275	20
vampirtrace/5.13	160	2
hpctoolkit/5.1.0	143	8

## Performance tools usage on Kraken

Library/ version	instances	users
perftools/5.2.0	3352	21
craypat/5.1.3/	2950	30
craypat/5.1.0/	1674	11
papi/4.1.2/	339	34
tau/2.20	315	38
fpmpi/1.1	295	11
papi/4.1.0.0.2	229	17
papi/3.6.2.2	200	8
papi/3.7.2	133	19
mpip/3.1.2	113	3

## Performance tools usage on Rosa

Library	instances	users
craypat	3465	32
papi	577	35
scalasca	163	7
tau	27	2

# Usage during executions

## Kraken

Applica-tion	instances	users	Application / version	instances	users
<b>namd</b>	368349	109	<b>namd/2.7/</b>	294547	19
<b>aprs</b>	192749	20	<b>namd/2.7b2</b>	16237	10
<b>amber</b>	71261	18	<b>namd/2.7b1-09</b>	7834	4
<b>hmc</b>	51541	10	<b>gromacs/4.5.3</b>	3162	12
<b>vasp</b>	17884	33	<b>namd/2.7b1/</b>	2576	3
<b>wrf</b>	20141	19	<b>amber/10</b>	1830	7
<b>espresso</b>	14597	20	<b>amber/11</b>	1081	4
<b>lammps</b>	7035	40	<b>namd/2.8</b>	1052	10
<b>gromacs</b>	6345	28	<b>cpmd/3.13.2</b>	1047	6
<b>cpmd</b>	1773	6	<b>q-espresso/4.2.1</b>	950	1

 MD applications

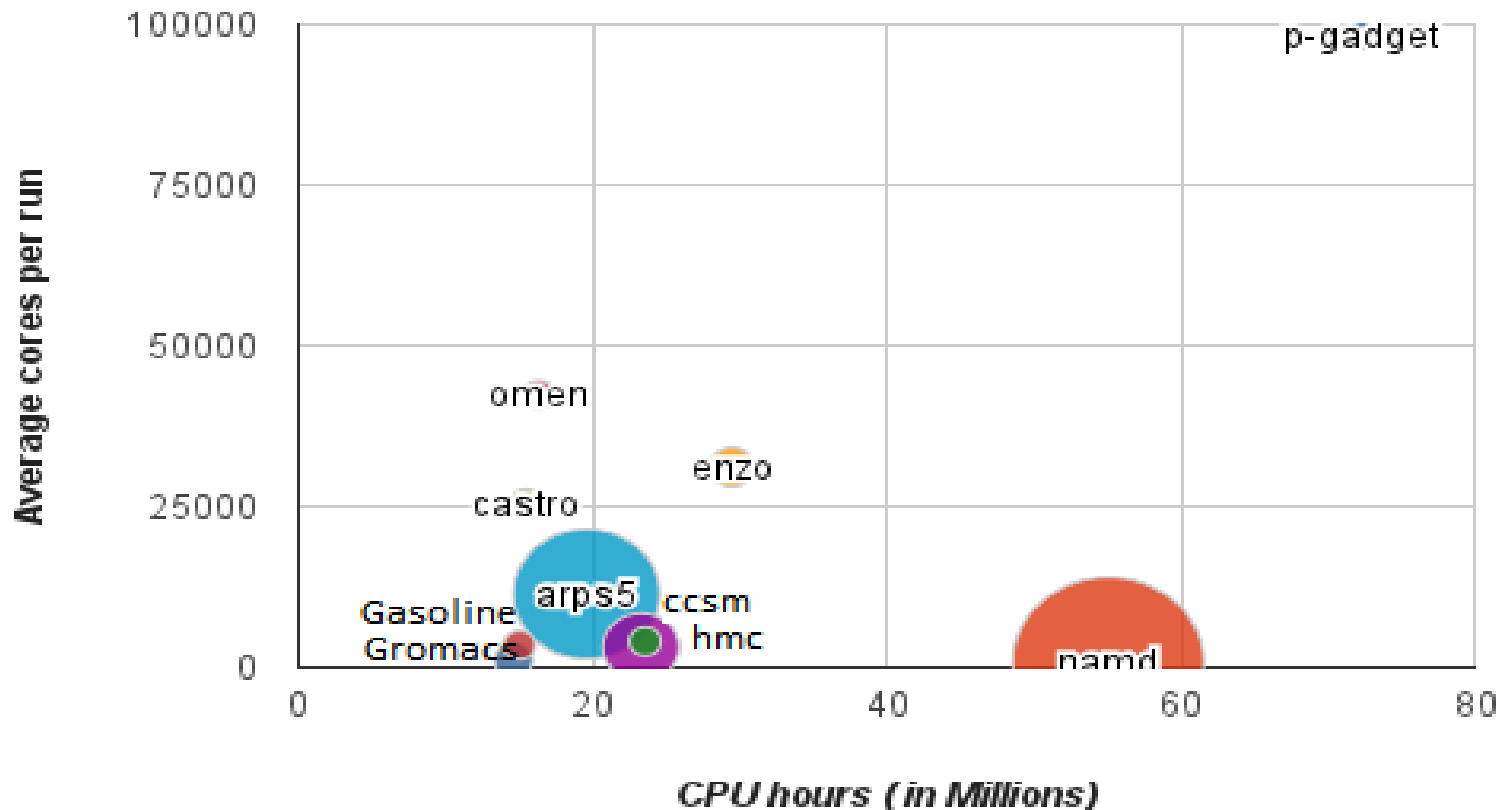
## Jaguar

## Rosa

Applica-tion	instances	users	Application / version	instances	users
<b>ior_bench</b>	496352	2	<b>vasp/4.6</b>	16333	8
<b>lammps</b>	105345	31	<b>lammps/9sep10</b>	2899	6
<b>esmf</b>	86480	9	<b>namd/2.7b1</b>	1400	4
<b>amber</b>	64725	4	<b>namd/2.6</b>	692	1
<b>vasp</b>	45533	26	<b>lammps/1jui11</b>	477	4
<b>nwpar</b>	33484	1	<b>spdcip/1.0.0</b>	121	17
<b>ccsm</b>	25557	61	<b>adios/1.3</b>	46	4
<b>espresso</b>	19605	20	<b>esmf/5.2.0</b>	25	3
<b>gromacs</b>	8443	9	<b>namd/2.7b4</b>	14	1
<b>namd</b>	5143	14	<b>gromacs/4.0.5/</b>	12	1

Applica-tion	instances	users	Application / version	instances	users
<b>namd</b>	74952	22	<b>namd/2.8</b>	29351	12
<b>int2lm1</b>	65274	8	<b>namd/2.7b4</b>	22160	6
<b>parfe</b>	52167	13	<b>espresso/4.2.1</b>	17296	7
<b>cp2k</b>	52010	37	<b>cp2k/21.11.2011</b>	11006	10
<b>siba</b>	37443	4	<b>cp2k/17.08.2010</b>	7838	5
<b>gromacs</b>	32573	13	<b>vasp/5.2</b>	2965	5
<b>echam</b>	28480	11	<b>vasp/4.6</b>	2053	3
<b>espresso</b>	21133	15	<b>espresso/4.3.2</b>	987	7
<b>dlpoly</b>	8535	4	<b>Cpmd/3.13</b>	680	4
<b>vasp</b>	5071	10	<b>espresso/4.1</b>	599	2

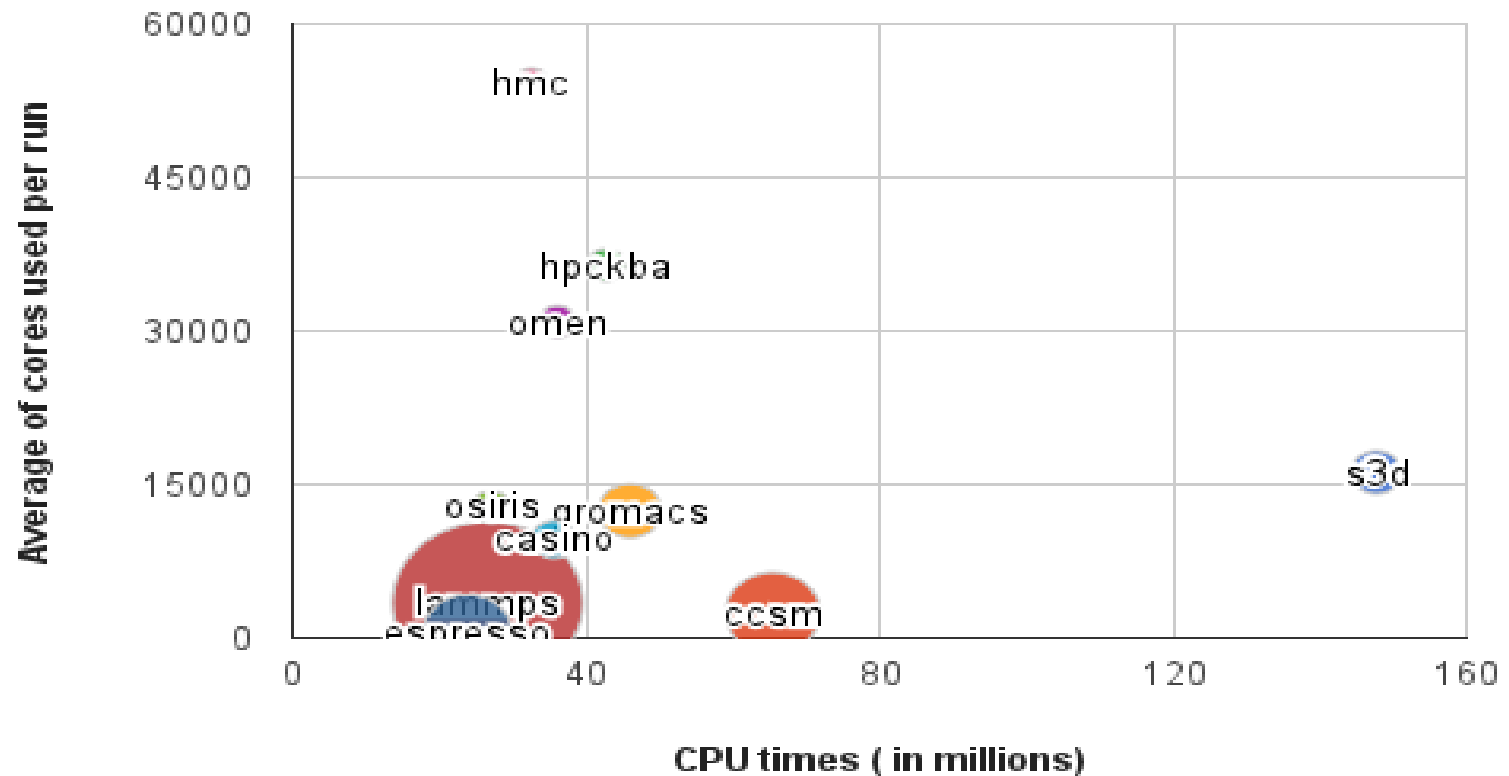
# Most CPU hours applications on Kraken



The top 10 most CPU consuming codes on Kraken in 2011, showing CPU hours consumed vs the average cores per run.



# Most CPU hours applications on Jaguar



# Contents

- Introduction and Motivations
- Overview on ALTD Tool
- Data Mining over 3 centers ( Kraken, Jaguar, Rosa)
  - Linkline
  - Execution
- **Conclusions and Future Work**

# Conclusions

- **ALTD tracks automatically and transparently library usage at compilation and at execution**
  - Wrapping the linker and the job launcher
  - In production on several Cray XT/XE machines at NICS and OLCF (ORNL), CSCS
  - **Alpha version is available if interested, please contact us !**
- **Track the most used libraries and it facilitates decisions for removing old/non-used libraries**
- **Data mining:**
  - **Usage at linking:**
    - Linear algebra, ( LibSci, FFTW, ACML, PETSc)
    - I/O ( HDF5, NetCDF)
    - Performance tools ( Craypat, PAPI, TAU)
  - **Usage at execution**
    - Molecular dynamic (NAMD, LAMMPS, GROMACS, ESPRESSO ), climate modeling(US),bone structures simulation ( EU) for the number of executions
    - Cosmology code for the CPU hours, ( Kraken), Combustion ( Jaguar)

# Conclusions

- **How ALTD has been useful to NICS?**
  - Detected bad behavior
  - Dropped Pathscale support
  - Upgrade to CLE3.1: prioritize tasks according to the high usage and choose the most used version along with the most recent one.
- **The results of our data mining have shown that a significant number of users are using their own libraries/executables even where there is a centrally installed application available:**
  - a user might require a non-standard version
  - prefer to use a version built by them
  - unaware that a centrally installed version exists. ( need more TEO )
  - Don't use optimized version

**→ Greengineering the present for the Future !**

# Future Work

- We need support to continue the efforts.
- The data mining has confirmed that there is extensive use made of numerical libraries (particularly LibSci, ACML, and PETSc).
  - envision adding functionality to track individual routines and correlate them back to “logical” libraries
    - assist library develops and vendors (and centers) to tune the most used functions for current and upcoming architectures, including multicore and accelerators.
- At present the data mining of ALTD is a manual process(python scripts) that generate a few simple SQL queries.
  - envision to provide tools that automate the querying process: support staff would thus be alerted immediately to cases where, for example, a user is running a code that is linked against a deprecated or buggy library

# Thank You !