

Linear Algebra libraries in Debian

$$\begin{pmatrix} 23 & 42 \\ 42 & 32 \end{pmatrix} + \begin{pmatrix} 0.3707945 & 0.1903269 \\ 0.2116117 & 0.5607954 \end{pmatrix}$$

Who I am ?

- Core developer of Scilab (daily job)
- Debian Developer
- Involved in Debian mainly in Science and Java aspects
- sylvestre.ledru@scilab.org / sylvestre@debian.org

What is the linear algebra ?

What is the linear algebra ?

- Linear algebra is a **branch of mathematics** concerned with the **study of vectors**, with families of vectors called vector spaces or linear spaces, and with functions that input one vector and output another, according to certain rules. These functions are called linear maps or linear transformations and are often **represented by matrices**. Linear algebra is central to **modern mathematics and its applications**.

Source: http://en.wikipedia.org/wiki/Linear_algebra

What is the linear algebra ?

- Widely used



Linear algebra libraries

Linear algebra libraries

BLAS (Basic Linear Algebra Subprograms)

- Developed in Fortran to manage matrix computations
- Very old (1979)
- API and ABI stable
- Provides vector operations (called L1), matrix-vector operations (called L2) and matrix-matrix operations (called L3)

Linear algebra libraries

LAPACK (Linear Algebra PACKage)

- Developed in Fortran too
- Built over BLAS
- Younger (1992)
- API and ABI stable
- Manage advanced operations on matrices for solving systems of linear equations, eigenvalue problems, and singular value decomposition...

Linear algebra libraries

But ...

- BLAS and LAPACK are now also *de facto* APIs
- To sum up:
BLAS is the API
REFBLAS is the reference implementation

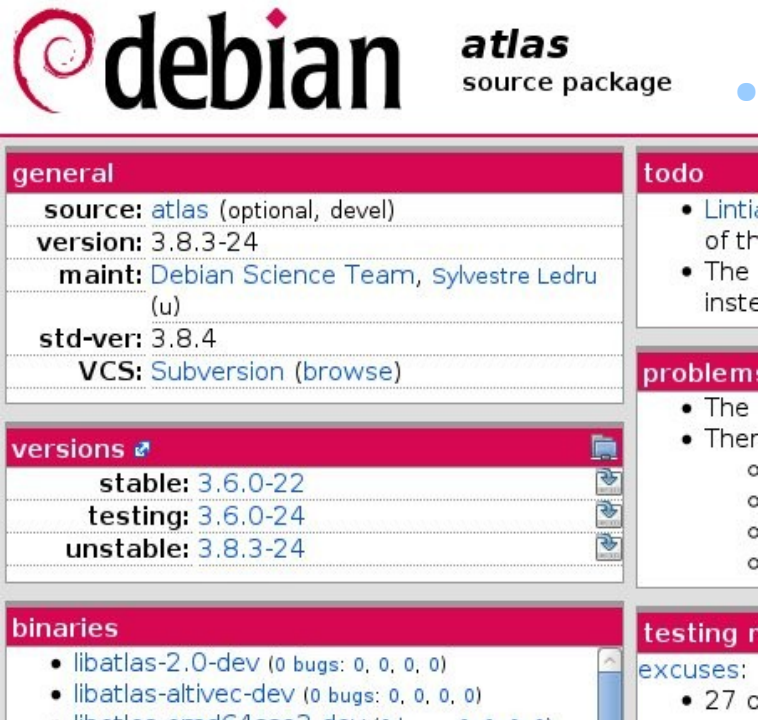
Linear algebra libraries

- REFBLAS and LAPACK performances are low.
Single threaded, no usage of CPU extensions, etc.
- Other implementations of the API:
 - Intel (MKL)
 - Sun (Sun performance library)
 - Apple (Velocity Engine)
 - AMD (ACML)
 - NVIDIA (CUBLAS)
 - ...
- Optimized for their architecture or operating systems
- Highly optimized libraries targeting HPC platforms

Status in Debian

Status in Debian

- Both REFBLAS and LAPACK in Debian Maintained by the Debian Science Team



The screenshot shows the Debian package page for 'atlas' source package. It includes sections for general information, versions, binaries, and a list of bugs.

general

- source: atlas (optional, devel)
- version: 3.8.3-24
- maint: Debian Science Team, Sylvestre Ledru (u)
- std-ver: 3.8.4
- VCS: Subversion (browse)

versions

- stable: 3.6.0-22
- testing: 3.6.0-24
- unstable: 3.8.3-24

binaries

- libatlas-2.0-dev (0 bugs: 0, 0, 0, 0)
- libatlas-altivec-dev (0 bugs: 0, 0, 0, 0)
- libatlas-cxx-dev (0 bugs: 0, 0, 0, 0)

todo

- Lintian
- of th
- The
- inste

problem

- The
- Ther

testing r

excuses:

- 27 c

- ATLAS (Automatically Tuned Linear Algebra Software) based on empirical techniques in order to provide portable performance. Available in Debian for various CPU extensions (SSE1, SSE2, SSE3, etc). Provides some great performances improvements in most cases. Provides all BLAS and some LAPACK functions

ATLAS: Status in Debian

But...

- Usage of optimized packages was hard in stable
- User had to play with `LD_LIBRARY_PATH`

ATLAS: Status in Debian

Proposal implemented in March 2010:

- Use the update-alternatives system for all BLAS and LAPACK implementations
- Similar approach to the MPI one

```
update-alternatives --config libblas.so.3gf
```

Example

There are 3 choices for the alternative libblas.so.3gf (providing /usr/lib/libblas.so.3gf).

Selection	Path	Priority	Status
* 0	/usr/lib/atlas-core2sse3/atlas/libblas.so.3gf	55	auto mode
1	/usr/lib/atlas-base/atlas/libblas.so.3gf	35	manual mode
2	/usr/lib/atlas-core2sse3/atlas/libblas.so.3gf	55	manual mode
3	/usr/lib/libblas/libblas.so.3gf	10	manual mode

ATLAS: Status in Debian

- Simplifies the switch between the implementations
- Will facilitate integration of other BLAS/LAPACK implementations
- For now, BLAS, LAPACK and ATLAS managed
- For more information:
<http://wiki.debian.org/DebianScience/LinearAlgebraLibraries>



But ?

The current ATLAS packaging is broken

Why ?

- Detection of CPU extensions is done once at build time
- The number of threads is computed at build time [1]
ie: I have a brand new 8 cores, amd64 atlas packages will expect all Debian users to have 8 cores too.
- Crashes at build time in some case (CPU extensions not available). 3 Rcs for now + numerical computing bugs.

[1] Atlas FAQ:

Can I vary the number of threads ATLAS uses dynamically?

No. The maximum number of threads to use is determined at compile time.

ATLAS proposal

- Drop optimized packages
- Update the libatlas3gf-base package description with custom build instructions
- Any comments ?

Other linear algebra libraries

Eigen - <http://eigen.tuxfamily.org/>

- C++ template based
- Good performances
- Easier and more intuitive to use
- Provide also a BLAS connector



Next steps

- Packaging of CUBLAS to NVIDIA GPU based (non-free)
- Packaging of GOTOBLAS (non-free)
- Integration of scalapack and the gsl (GNU scientific library) into the update-alternatives system

Help is welcome.

Questions, comments ?