

A Collaborative guide to ATLAS Development

R. Clint Whaley ^{*}Peter Soendergaard [†]

February 2, 2007

Abstract

This paper exists to get you started if you want to do some ATLAS development. The hope is that as new developers work on ATLAS, they will add to this note, so that this information grows with time.

^{*}`rwhaley@users.sourceforge.net`

[†]`soender@users.sourceforge.net`

Contents

1	Introduction	3
2	Adding to this note	3
3	Getting the ATLAS code through CVS	4
3.1	Background on ATLAS code base	4
3.2	Getting the ATLAS CVS tree and the working ATLAS directory	4
3.3	Basefile/extracted file interaction	5
4	Getting CVS write access/Becoming an ATLAS SourceForge developer	5
4.1	ATLAS developer guidelines	5
4.2	Getting CVS write access	6
4.2.1	I don't agree to the guidelines, what now?	6
4.2.2	How about being a developer and not using CVS for writes?	6
4.3	Setting up for CVS write access	7
4.4	Further info	7
5	Adding a LAPACK routine to ATLAS	8
5.1	Row-major LAPACK routines	8
5.2	Outline of Steps	9
5.3	Create and debug tester	9
5.3.1	Writing ATLAS/src/testing C-to-f77 wrapper	10
5.3.2	Getting your routines extracted	10
5.4	Create and debug ATLAS internal routines	10
5.5	Add C and F77 interface routines	10
5.6	Update the LAPACK quick reference guides	11
6	Architectural defaults	11
6.1	Rambling on about architectural defaults	11
6.2	Making your own architectural defaults	12
6.3	Getting ATLAS to use your shiny new defaults	12
7	Sanity testing for an ATLAS install	12
7.1	Invoking the sanity tests	13
7.2	Understanding the sanity test output	13
7.3	Finding the context of the error	14
7.4	Tracking down an error in the BLAS interface testers	14
7.5	Tracking down an error in the bin/ testers	15
8	Antoine's testing scripts	16
8.1	Setting up and installing the tester	16
8.2	Getting a good BLASlib	16
8.3	Running the tester	17
8.4	Finding errors	18
8.5	Tracking down errors in the bin/ testers	18

9	Finding a good NB for GEMM	18
9.1	Better kernel timing does not always yield faster GEMM	19
9.2	Large NB means more time in cleanup	19
9.3	Large NB means more time in unblocked application code	19
9.4	Finding a good NB	20
10	Information on atlconf	22
10.1	Weaknesses in <code>configure</code>	22
10.2	Weaknesses in <code>spew/config</code>	23
10.3	Probe Overview	23
10.4	Architectural Probes	25
10.5	Notes on <code>configure</code>	26
10.6	Compiler Setup and Handling in ATLAS Config	26

1 Introduction

This note exists to get you started if you want to write **new** code for ATLAS, or if you want to modify ATLAS source. It is *not* for kernel contribution, which is what most people do when contributing to ATLAS. Kernel contribution is much simpler, and is explained in `ATLAS/doc/atlas_contrib.ps` or in html format at:

http://math-atlas.sourceforge.net/atlas_contrib

So what is the difference between development and kernel contribution? In kernel contribution, you write a kernel to be used by ATLAS, using the provided ATLAS testers and timers to verify it, and when you are satisfied with its performance and reliability, you submit your kernel to the ATLAS team, and they accept it or not, and you are done.

Doing actual core development is quite a bit more complex. Probably the biggest change is that you will need to write your own tester and timer for your new code. No code will be accepted into the ATLAS code base without a tester which can be used to verify it. Since writing a decent tester is usually at least as hard as writing the code it tests, and is always a whole lot less enjoyable, the author must bear the pain of producing it along with the pride of producing the code. As a developer, you will be responsible for testing such new code on several platforms as well.

If you are instead hoping to modify some of the existing code base, remember that for non-kernel operations, portability and robustness *must* be the primary goal. There are many sections of ATLAS that we know to be second rate on a certain platform, but we also know that it works on the twenty or so architectures that ATLAS is routinely compiled for, so we leave it that way. This means that when a modification is made to a previously existing routine, the modifying author must have good evidence that the new code is as portable as the old. In short, the barrier to replacing tested code is high.

It is possible that users want access to the CVS repository even though they do not plan on doing development, mainly I'd guess so they can get access to the newest stuff without waiting for developer releases. Also, kernel contributors who make subsequent changes to their routines can speed up their adoption by submitting them in a format ready for CVS check-in. =====

2 Adding to this note

This note is included in the `AtlasBase/TeXDoc` of the CVS repository, and anyone can submit a patch against it giving additional information. As the founder of ATLAS, I have written a seed of a document explaining how to get access to the code. It is my hope that other developers will add important information that they discover as they go, so that this doc will grow over time, getting much information that I probably take so much for granted that I would never think to document.

Therefore, new sections are welcome, and probably a FAQ appendix would be a good idea. As people contribute, their names will be added to the author list.

3 Getting the ATLAS code through CVS

3.1 Background on ATLAS code base

ATLAS was not originally developed in CVS. ATLAS was developed using a programming tool called *extract*, which means ATLAS is actually maintained in something called base-files. If you think of regular development being access by value, and CVS using a level of indirection, CVS on basefiles gives you two such levels of indirection. So, if you want to be able to use CVS check-ins, you will need to learn at least the basics of *extract*. Details on *extract* can be found at:

<http://www.cs.utk.edu/~rwhaley/extract/>

Note that if you just want read-only access, you will need to install *extract* so that you can get at the files, but you will not need to learn anything about it.

3.2 Getting the ATLAS CVS tree and the working ATLAS directory

Here are the steps to get a local ATLAS CVS tree:

1. **Install *extract*:** Download and install *extract* as described here:

<http://www.cs.utk.edu/~rwhaley/EXTRACT/faq.html#install>

2. **Checkout the ATLAS CVS tree:** From a CVS-capable machine connected to the Internet, go to the directory where you want the ATLAS *basefiles* to be, and issue:

```
cvs -z3 -d:pserver:anonymous@math-atlas.cvs.sourceforge.net:/cvsroot/math-atl
as checkout AtlasBase
```

3. **Update *topd*:** Edit the created *AtlasBase/make.base*, and change the definition of the *topd* macro to match where you have put the *AtlasBase* directory. For example, if I issued the CVS command in */home/rwhaley/work*, I change line 46 from something like:

```
@define topd @/home/rwhaley/Base/AtlasBase@
```

to:

```
@define topd @/home/rwhaley/work/AtlasBase@
```

4. **Create ATLAS working directory and *extract* files:** Go to the directory where you want the working copy of the ATLAS/ directory to reside. In there, issue these two commands:

```
extract -b <your topd>/make.base -o Makefile rout=Make.atldir
make
```

3.3 Basefile/extracted file interaction

You'll have to scope the extract page for any kind of real feel for how this works, but some atlas-specific details are in order here. In each subdirectory of the `ATLAS/` tree, you will find a file `Make.ext`. If you type `make -f Make.ext`, this makefile will extract all new files in this subdirectory for which the basefile is newer.

So, what usually happens is you are messing with something, and you do it in the `ATLAS/` directory. When you are confident in your change, you put it into the appropriate basefile in the `AtlasBase/` directory (note that examining `Make.ext` will show you what basefile a given extracted file comes from), and you then re-extract over your working copy with the above command.

4 Getting CVS write access/Becoming an ATLAS SourceForge developer

In a colossal case of lunacy, the developers of SourceForge decided to allow anyone associated with a SourceForge project full write access to all CVS files. The permissions that can be granted or taken away are the right to administer lists, changes html docs, do file releases, etc. To me, this is a little like having a bank that encrypts your name and password with the best available means, but leaves your money out in open bins on the front lawn.

Normally, I would say the more the merrier in terms of adding people as SourceForge developers. I'd be happy to see hundreds of people associated with the project. On the other hand, I'd be a little scared with having hundreds of people I've never met have full write access to a project as detailed and delicate as ATLAS. CVS has marvelous rollback abilities, but I'm afraid as of right now anyway, I don't have marvelous CVS abilities, and so I tend to err on the side of timidity.

What all this whining is coming down to is, you have to first show me the money before I'll add you as a developer :). When you first begin hacking ATLAS, the method to use is to submit patches or code to the list, and if your submission is incorporated into ATLAS, we can then get you added as a SourceForge developer.

4.1 ATLAS developer guidelines

We're all still learning, but here are the rough guidelines that you'll need to be OK with to be added as an ATLAS SourceForge developer:

1. **You only directly commit to basefiles you originate.** Assuming you have written something for ATLAS before, you should have your own source to maintain. For instance, if you add some new kernels, we'll create you a subdirectory under `AtlasBase/kernel`. For some development, though, you will need to modify basefiles that you did not originate. For instance, if you wrote a new lapack routine, you would minimally need to modify a basefile originated by me that creates the lapack header files. For these mods, you should ideally send the basefile maintainer a patch against the basefile, which the basefile maintainer will vet and apply. If that is impractical, you will need to agree with the maintainer on a different strategy. This rule is very important, and I will tend to remove someone from the group immediately if it is violated.

2. **Always leave the main CVS branch in a working, tested state.** As completely unfair as it is, I (Clint) am the only guy allowed to break the main CVS branch. This means that if you have serious development you'd like to do under CVS, you need to do a developer branch until everything is working. If your mods are all to your own files, you will simply merge the developer branch back into the main once the new stuff is tested and working. If they involve other routines, then you'll need to have the person responsible for the modified codes help with the merge. So what can you apply directly to the main trunk? For instance, if you were maintaining a kernel, and you improved it, then tested (building the full gemm, running the testers, etc), and everything is OK, you could check that in **after** the testing and development process was done.
3. **You only submit on the bug tracker for your own code.** The ATLAS bug tracker serves roughly the same service for developer releases the errata does for stable: a place for all confirmed bugs. You can confirm bugs in your code, but the maintainer of the problematic code should confirm bugs in code he/she is responsible for. So, if you find what you think is a bug in someone else's code, you submit it as a support request, and the code maintainer should scope it there. If it is indeed a bug, the maintainer should move it to the bug tracker, where it will remain until fixed by a subsequent developer release. Often, these "bugs" may turn out to be misunderstandings or bad installs, and thus never get moved to the bug list.
4. **When in doubt, ask.** If you want to do something that you think might effect other people, ask them first. If you are not sure if you should do something, ask. If the question is general, post to the developer list. If it is delicate or personal in some way, you can send to me directly at `rwhaley@users.sourceforge.net`.

4.2 Getting CVS write access

If the developer guidelines seem reasonable to you, and you have code that is already been accepted into ATLAS, send mail to `rwhaley@users.sourceforge.net`, saying you want CVS write access, and agree to the developer guidelines. Just to make sure you've read them, I'd like the e-mail to specifically say:

```
I will not do CVS check-ins on files that I am not the maintainer
for, and I'll keep the main branch in working order.
```

4.2.1 I don't agree to the guidelines, what now?

Send it to the list, or to me if it is delicate. We ought to be able to find a way to make things work.

4.2.2 How about being a developer and not using CVS for writes?

Say you want to write a kernel for the AltiVec, but you don't presently have access to a G4. SourceForge has a compile farm that contains a G4, but you don't have compile farm access until you are a member of a SourceForge group. In this case, I'm happy to add you as a developer with your agreement that you won't do CVS writes until we agree otherwise.

4.3 Setting up for CVS write access

After we agree that you will be an ATLAS SourceForge developer, you need to go to www.sourceforge.net, and sign up as a SourceForge user. You should see a sidebar with:

Status:

Not logged in

Login via SSL

New User via SSL

You want to take the `New user via SSL` link. This gets you a SourceForge user name, which you need to send to me at `rwhaley@users.sourceforge.net`. I then have to add you as an ATLAS developer.

You need to change your CVS access from anonymous/read to developer/write. It helps if you set your `CVS_RSH` and `CVSRROOT` environment variables appropriately. `CVS_RSH` should be `ssh`. My `CVSRROOT` is set to

```
:ext:rwhaley@math-atlas.cvs.sourceforge.net:/cvsroot/math-atlas
```

Replacing my name (look for `rwhaley` in the above) with yours should get you what you need.

CVS check in/out will still not work correctly at this date. Now, do:

```
ssh math-atlas.cvs.sourceforge.net
```

and enter your user name and password. As soon as you get a prompt, logout. This process creates some needed files.

Finally, if you have an ATLAS tree created by anonymous CVS access, the easiest thing is to delete it and recheck out as yourself. CVS creates some files on the first checkout saying what kind of access you are using, and these will still show as anonymous, despite your `CVSRROOT` (this is true for branching as well, so watch out). If you want to scope the CVS commands, I'm sure you can switch over without deleting, but removing and rechecking out is how I've seen this problem successfully fixed.

4.4 Further info

As I said before, I'm not much of a CVS guru. I have found

```
http://www.loria.fr/~molli/cvs/doc/cvs\_toc.html
```

to be a helpful in doing CVS stuff.

For SourceForge, everything I know about it came from:

```
https://sourceforge.net/docman/?group\_id=1
```

The companion piece of this guide is the atlas contributor guide, found in `ATLAS/doc/atlas_contrib.ps` of the tarfile. You really need to know that before doing much with this guide.

Extract is explained at:

```
http://www.cs.utk.edu/~rwhaley/extract/
```


5 Adding a LAPACK routine to ATLAS

It is probably not practical that ATLAS will provide a complete LAPACK API (as it does with the BLAS) in the foreseeable future, both due to the algorithmic complexity of some of the operations, and to the sheer number of routines in LAPACK. It must be understood that adding routines adds to the inertia and maintenance costs of the package, and this additional burden must be offset by real advantage for the user.

ATLAS has so far only added LAPACK routines to ATLAS when we can make a performance-enhancing algorithm modification. For instance, we added the LU and Cholesky factorizations because we used the recursive formulations of these routines, which provides for better performance on pretty much any cache-based architecture.

However, when we have added such routines, we usually add the correlated routines even when a performance advantage is not supplied. For instance, upon adding GETRF support, we also added GETRS and GESV. As far as column-major routines go, we supply no better algorithm for GETRS or GESV than LAPACK. However, since these routines are very simple, and GETRF is very often used with them, we added them along with GETRF. The idea here is that their maintenance costs are not heavy, and real advantage is given to the user in that we have sped up GETRF, and if the factor and solve are all he needs, ATLAS will supply a complete solution.

The column-major comment points out another reason to add a routine to ATLAS: ATLAS supplies the only performance-aware row-major LAPACK implementation that I am aware of (I'm sure there are some, I just don't know of any that aren't simply using the col-major stuff, and thus performing terribly). It is possible that someone would want to add an LAPACK routine to ATLAS simply because they need a row-major version, and someone being motivated enough to write it would probably be ample justification to add the routine to the ATLAS tarfile.

5.1 Row-major LAPACK routines

So far, we have accepted no routines that do not also include a row-major equivalent, both for BLAS and for LAPACK. We hope to continue this. There are as yet only a few users of the row-major LAPACK/BLAS that I am aware of, but I believe that this is a chicken/egg problem.

Some people insist on using row-major arrays in C, but if they have access to a BLAS/LAPACK that supports it, they find the performance is no better than what they get with simple loops, or that it is calling the col-major in a naive way, and cutting the problems size they can solve in half by copying. Therefore, people with row-major bias don't call the stuff 'cause it doesn't help them, and the problem continues.

It is my belief, therefore, that good-quality row-major stuff must be produced before significant demand will appear. If I'm wrong, I guess we'll someday drop support for row-major, but I don't think this will be the case over a long enough time line.

Therefore, despite it being a hassle, having a good quality row-major implementation is critical for getting an LAPACK routine into ATLAS. For many routines, since we have row-major BLAS, the algorithm stays the same, and only some pointer arithmetic need be changed.

Other routines in LAPACK (GETRF is one) have a built in algorithmic bias towards column-major (in GETRF, this is doing row-pivoting), and another algorithm with the same

stability and usage characteristics should be employed for row-major (eg., column-pivoting, for GETRF).

5.2 Outline of Steps

Here are the general steps to use when adding an LAPACK routine to ATLAS:

1. Create and debug tester using LAPACK
 - Update extractors
 - Update Makefiles
2. Write and test ATLAS internal routines using above tester
 - Update extractors
 - Update Makefiles
3. Update `atlas_clapack.h` to include your new routines
4. Create C and F77 interfaces to your routine
 - Update extractors
 - Update Makefiles
5. Update `clapack.h`
6. Update the LAPACK quick reference guides.

5.3 Create and debug tester

The first step in adding a new routine to ATLAS is to create a tester (and timer) which can be used to verify the correctness of your code. More than half of the challenge is getting the tester right; with a good tester/timer, the code usually comes fairly easily.

Your tester will go in `ATLAS/bin` when extracted; you can examine some of the testers available there to get an idea of what you should do (eg., look at `ATLAS/bin/[lu/llt/slv/trtri/uum]tst.c`). All of these routines come from the basefile `AtlasBase/Clint/atlas-tlp.base`, which is what you should submit your patch against, unless you want to create your own, separate basefile.

After your tester is written, it's column-major components can be tested against LAPACK by using the `make <rout>tstF` target in `ATLAS/bin/<arch>`. You can even test the row-major components by having the F77 interface transpose the matrices on input, and back on output. See `ATLAS/bin/uumtst.c` for an example of this for square matrices.

As part of your debugging of the tester, be sure that it not only agrees that LAPACK produces the right answer, but truly detects errors as well. For instance, manually overwrite an entry, both in the matrix and in the padding (in separate tests), and make sure it is caught by the tester.

5.3.1 Writing ATLAS/src/testing C-to-f77 wrapper

You first need a way for your tester, written in C, to call the LAPACK routine, written in F77. All such language translation routines are kept in `ATLAS/src/testing`, and come from the basefile `ATLAS/Clint/atlas-ilp.base`. This wrapper is trivial, though some of the integer/string stuff is not obvious. Steal the code from the other examples.

5.3.2 Getting your routines extracted

Now you need to get your files to appear in the right subdirectories, so you need an entry in the appropriate `Make.ext`. All the `Make.ext` files comes from `AtlasBase/make.base`, so find the rout for your directory in this file (for examples the line saying `@ROUT ATLAS/interfaces/lapack/F77/src/`) and add your routine name to the line containing the name of all the other routines.

So, now you are in your working directory (say `ATLAS/src/lapack`), and you type `make -f Make.ext`, and nothing happens, no new files show up. This is because you need to re-extract your `Make.ext` file. This can of course be done by removing your whole ATLAS tree and reinstalling, but less brutally you can “just” use something like this: `extract -b /home/soender/AtlasBase/make.base -o Make.ext rout=ATLAS/src/lapack -langM`. The `-langM` switch is required for `extract` to properly handle makefiles, so you cannot skip it.

This is the basic procedure for this sort of stuff. When you need a makefile in an ARCH subdirectory, the appropriate makefile is copied by `Make.top` from the `ATLAS/makes/` directory. Check `Make.ext` to see which basefile they come from, and add your routine name among the names of the other routines.

Remember to update the Makefiles for both `ATLAS/bin` and `ATLAS/src/testing`, and to get these makefiles into the appropriate subdirs. In order to extract new makefiles, and get them put into the appropriate subdirs, I typically do something like:

```
pushd ~/TEST/ATLAS/makes/ ; make -f Make.ext ; cd .. ; \  
make refresh arch=Linux_PIIISSE1 ; popd
```

(replace the path and arch appropriately, obviously).

5.4 Create and debug ATLAS internal routines

The internal LAPACK routines are kept in `AtlasBase/Clint/atlas-lp.base`. Add your routine here, and update `ATLAS/src/lapack's Make.ext` and `Makefile` appropriately to build your routine.

You will add your routine in `atlas-lp.base` with an additional `@ROUT` keyline, but also do not forget to update the include file `atlas_lapack.h` at the bottom of the file as well. You will need to add your routine to the prototype part, as well as to the macro renaming part. Examine the basefile for details.

Once it is extracting, use your LAPACK-debugged tester to debug your code.

5.5 Add C and F77 interface routines

We do this step last, because we don't want to add API routines until the code is working. Having debugged and made sure the code is faster than LAPACK, we're now ready to make it available to the user via the advertised APIs. The extracted API files are kept in subdirectories under `ATLAS/interfaces/lapack`.

The F77 interfaces are kept in `AtlasBase/Clint/atlas-fint.base`. Look at the existing examples and notice how `extract` generates all four precision from the same routine, if you use the `extract` macros. All the code for this interface can be ripped from LAPACK and adapted. Note that you will usually need to examine both complex and real versions of the original LAPACK routine, to find any differences in interface/testing and comments. You will also need to remove unneeded `EXTERNAL` declaration, etc.

This interface does the parameter checking, and converts any FORTRAN string arguments to some predefined integer values, and then call the `ATLf77wrap` interface. Scope any of the existing routines for details on this.

The C interfaces are easy to write, since they should just check the input arguments, and then call the ATLAS routine. The codes are stored in `atlas-clp.base`. Check it out for lots of examples.

5.6 Update the LAPACK quick reference guides

The ATLAS user API is defined in the quick reference guides under `AtlasBase/TexDoc`. Right now, the supported LAPACK API is small enough to fit both C and F77 interfaces on one card (single 2-sided landscape page), but eventually it will be split in two, as with the BLAS quick reference cards. Either use the Makefile to do it, or remember to manually throw the `-tlandscape` flag to `dvips`, and the `-paper a4r` flag to `xdvi`.

6 Architectural defaults

ATLAS's architectural defaults are simply a record of the results of a previously run ATLAS search. They exist for a couple of reasons:

1. Using architectural defaults, install times are reduced to almost bearable levels
2. Because the search is empirical, installs can go wrong if unmonitored. Architectural defaults given out in the standard tarfile have at least passed the laugh test

6.1 Rambling on about architectural defaults

One FAQ for architectural defaults is why any timings are necessary when using them. The standard architectural defaults only rarely describe everything discovered by a search, but rather give only those data that we feel sure will not vary a great deal. For instance, for many machines, the kernels to use, etc., are fully specified, but `CacheEdge` is not. `CacheEdge` varies depending on your L2 cache size, which varies depending on architecture revision, so it is not specified, allowing it to tune itself for this variable parameter, while still skipping the search over less variable things (eg., if the L1 cache or FPU units change, this is usually a new architecture, not a revision of an old).

That's the theoretical reason why they shouldn't cover all discovered items. However, ATLAS presently times the kernels in order to be able to produce a comprehensive SUMMARY.LOG, and these timings *could* be skipped, assuming this functionality were added to the atlas install process.

There are some weaknesses of architectural defaults. One of the main ones is how they can go out of date, and cause slowdown. One big way this can happen is with compiler

changes. For instance, gcc 3.0 produces completely different (and inferior) x86 code than the 2.x series. On Athlons, performance is cut almost in half when the gcc 2.x defaults are compiled with gcc 3.0. Some of this performance loss is irrecoverable, due to 3.0's shoddy handling of the register stack. Other parts, however, can be made up by having ATLAS adapt itself to the new compiler, and architectural defaults prevent this from happening.

6.2 Making your own architectural defaults

This section describes how to create architectural defaults as of ATLAS release 3.7.12. For older releases, the process is similar, but not quite the same.

1. Get an install, correct in all details, that you want to immortalize.
2. cd to your <OBJdir>/ARCHS directory
3. Type `make ArchNew`

This will copy the search result output files into a directory <OBJdir>/ARCHS/<MACH>/, with appropriate subdirs under that. You can then go into these guys and delete files you don't want to be part of the defaults (eg., `atlas_cacheedge.h`, etc).

Now, to save these defaults to a transportable format, you can have the makefile create the tarfile for you by:

```
make tarfile
```

6.3 Getting ATLAS to use your shiny new defaults

Pretty easy:

1. Run configure, creating `Make.inc`, but do not start the install.
2. Take the tarfile you created, and copy it under `ATLAS/CONFIG/ARCHS` source directory.
3. Edit your `Make.inc` and make sure the `ARCH` macro matches the base name of your tarfile (eg., `P4ESSE3`), and that the `INSTFLAGS` macro has the flags `-a 1` (do use arch defs).
4. Continue the install as normal (eg. `make build`).

7 Sanity testing for an ATLAS install

From ATLAS3.3.8 forward, ATLAS has had a "sanity test", which just does some quick testing in order to ensure that there are no obvious problems with the installed ATLAS libraries. It runs all of the standard BLAS interface testers, with the default input files, and it then runs a few fixed cases of ATLAS's lapack tester routines (eg., `ATLAS/bin/invstst.c`, etc). The advantage of these lapack testers is that they depend on many of the BLAS as well as the lapack routines, so you get a lot of testing for a minor amount of time. The sanity checks do not require any non-ATLAS libraries for testing, so the only dependence that a user who has installed ATLAS may not be able to satisfy is the need for a Fortran77 compiler, which is required for the BLAS interface testers. As of ATLAS3.7.12, ATLAS also can run a reduced set of tests for users who do not have a fortran compiler.

7.1 Invoking the sanity tests

These tests are invoked from your install directory by:

```
make test
```

If you are using threads, you will want to run the same tests for threading via:

```
make pttest
```

The C-only testing subset can be invoked by prepending “C_” to the test name (eg., “make C_test”).

7.2 Understanding the sanity test output

Once you fire off this tester, you’ll see a lot of compilation going on. All compilation is done up front, and then the testers are run at the end. All tester output is dumped to some files (we’ll see specifics in a bit), which are then automatically grepped for errors at the end of the run. It is the results of this grep that the user will see. For example, here’s the output from a run on my Athlon running Linux:

```
dudley.home.net. make test
...
... bunch of compilation ...
...
DONE BUILDING TESTERS, RUNNING:
SCOPING FOR FAILURES IN BIN TESTS:
fgrep -e fault -e FAULT -e error -e ERROR -e fail -e FAIL \
    bin/Linux_ATHLON/sanity.out
8 cases: 8 passed, 0 skipped, 0 failed
4 cases: 4 passed, 0 skipped, 0 failed
8 cases: 8 passed, 0 skipped, 0 failed
4 cases: 4 passed, 0 skipped, 0 failed
8 cases: 8 passed, 0 skipped, 0 failed
4 cases: 4 passed, 0 skipped, 0 failed
8 cases: 8 passed, 0 skipped, 0 failed
4 cases: 4 passed, 0 skipped, 0 failed
DONE
SCOPING FOR FAILURES IN CBLAS TESTS:
fgrep -e fault -e FAULT -e error -e ERROR -e fail -e FAIL \
    interfaces/blas/C/testing/Linux_ATHLON/sanity.out | \
    fgrep -v PASSED
make[1]: [sanity_test] Error 1 (ignored)
DONE
SCOPING FOR FAILURES IN F77BLAS TESTS:
fgrep -e fault -e FAULT -e error -e ERROR -e fail -e FAIL \
    interfaces/blas/F77/testing/Linux_ATHLON/sanity.out | \
    fgrep -v PASSED
make[1]: [sanity_test] Error 1 (ignored)
DONE
```

So, in the LAPACK testers we see no failures (all tests show 0 **failed**), and we have no output from the BLAS testers, which is what we want. Notice the lines like:

```
make[1]: [sanity_test] Error 1 (ignored)
```

This is due to fgrep's behavior, and does not indicate an error. If fgrep does not find any pattern matches, it returns a 1, 0 on match. Therefore, since we are grepping for error, getting an "error condition" of 1 is what we hope for.

7.3 Finding the context of the error

If the sanity test output shows errors, the next step is to track down where they are coming from. You can see in the output the files that are being searched for errors. They are:

```
bin/sanity.out
interfaces/blas/C/testing/sanity.out
interfaces/blas/F77/testing/sanity.out
```

The threaded sanity test uses the same filenames with **pt** prefixed.

The first thing to notice is which of these tests are showing errors. The testers in **bin** are higher level than those in the **interfaces** directories, so if you get errors in both, track down and fix the interface errors first, as they may be causing the lapack errors. If both **C** and **F77** BLAS interfaces are showing errors, I always scope and fix the Fortran77 stuff first, since Fortran is simpler (no RowMajor case to handle). Only if an error only shows up in **C** testing do I scope that output instead of the Fortran77.

The grepped error message probably gives you no idea what actually went wrong (it may show something as simple as:

```
FAIL
```

for instance), so you must go look at the **sanity.out** in question. For instance, you might need to scope **interfaces/blas/F77/testing/sanity.out**. You do a search for whatever alerted you to the problem (eg., **FAIL**), and you see by the surrounding context what tester failed.

7.4 Tracking down an error in the BLAS interface testers

The BLAS testers are split by BLAS Level (1, 2 or 3) and precision/type (s,d,c,z). The basic names of the tester executables are

```
x<pre>blat<lvl>
x<pre>cblat<lvl>
```

for Fortran77 and C, respectively. The Level 1 testers (**x[s,d,c,z]blat1**) test certain fixed cases, and thus take no input file. So if the error is in them, you simply run the executable with no args in order to reproduce the failure.

The Level 2 and 3 testers allow a user to specify what tests should be run, via an input file. The standard input files that ATLAS runs with are:

```
<pre>blat<lvl>.dat
c_<pre>blat<lvl>.dat
```

respectively. The format of these input files is pretty self explanatory, and more explanation can be found at:

www.netlib.org/blas/faq.html

To run the tester with these files, you redirect them into the tester. For instance, to run the double precision Level 2 tester with the default input file, you'd issue:

```
./xdbl2 < ~/ATLAS/interfaces/blas/F77/testing/dblat2.dat
```

You should be aware that only the first error report in a run is accurate: one error can cause a cascade of spurious error reports, all of which may go away by fixing the first reported problem. So, it is important to find and fix the errors in sequence.

I usually copy the input file in question to a new file that I can hack on (for instance, if the error was in the double precision Level 2, I might issue:

```
cp ~/ATLAS/interfaces/blas/F77/testing/dblat2.dat bad.dat
```

I then repeatedly run the routine and simplify the input file until I have found the smallest, simplest input that displays the error.

The next step is to rule out tester error. The way I usually do this is to demonstrate that the error goes away by linking to the Fortran77 reference BLAS rather than ATLAS (you can only do this for errors in the F77 interface, obviously). I usually just do it by hand, i.e., for the same example again, I'd do:

```
f77 -o xtst dblat2.o /home/rwhaley/lib/libfblas.a
```

If the ATLAS-linked code has the error, and this one does not, it is a strong indication that the error is in ATLAS. If the F77 BLAS are shown to be in error, it is usually a compiler error, and can be fixed by turning down (or off) the optimization used to compile the tester.

Now you should have confirmed the tester is working properly, and that the error is in a specific routine (let us say DNRM2 as an example). As a quick proof that DNRM2 is indeed the problem, you can link explicitly to the F77 version of DNRM2, and to ATLAS for everything else (see Section 8.2 for hints on how to do this). If this still shows the error, you are confident that ATLAS's DNRM2 is indeed causing the problem, and you should either track it down, or report it (depending on your level of expertise).

7.5 Tracking down an error in the bin/ testers

The sanity tests only run the LAPACK testers in this directory. The LAPACK routines depend on the BLAS, so ignore errors in lapack testers until all the BLAS pass with no error. If you have errors in LAPACK but the BLAS pass all tests, then you have to hunt for the error in the LAPACK routines.

First, rule out that it's not a problem in the BLAS that is just not showing up in the BLAS testing. Get yourself a reference BLAS library, as explained in Section 8.2. Then, set your `Make.inc`'s `BLASlib` macro to point to the created reference BLAS library. Then, you need to compile a library that uses ATLAS's lapack routines, but the reference BLAS. This can be done by compiling the same executable name with `_sys` suffixed. For instance, if you were running the LU tester, `xdlutst`, you would say `make xdlutst_sys`, and then run this executable with the same input.

If the error goes away, then the error is really in the ATLAS BLAS somewhere. I then usually look at the LAPACK routine and tester in question to find out what its BLAS dependencies are, and manually link in the reference BLAS object files until I find the exact BLAS causing the problem. Usually once you know what routine causes the prob, you can reproduce the error with the BLAS tester (i.e. you need a IDAMAX call with N=12, incX=82).

If the error still persists using ATLAS's LAPACK and the Fortran77 BLAS, the next trick is to do LAPACK just like the BLAS: download and compile the F77 LAPACK from netlib (www.netlib.org/lapack/lapack.tgz). You then set your `Make.ARCH`'s `FLAPACKlib` to point to your Fortran77 lapack library. You then suffix the base executable name with `F.sys` (eg., for LU again, you would do `make xdlutstF.sys`), and you will get a tester linked against the Fortran77 BLAS *and* LAPACK. If this also shows to be in error, there is an error in the tester, or in the compiler. Try turning down compiler optimization to rule in or out compiler errors.

8 Antoine's testing scripts

I have not updated these guys to use the 3.7.12 way of doing things yet!

Before a stable release, we always do as much testing as possible. The 900 pound gorilla of testers is Antoine's tester scripts. This tester runs for several days, and does a great number of both fixed and random tests, and if it completes with no errors, you have a pretty good idea that the code is fairly solid. This tester takes a lot of time and energy to get going, so not everyone will want to apply it. Even the casual user ought to run the sanity testing as a matter of course, and that should always be ran and passed first. Also, much of the methodology for understanding output, tracking down problems, etc, is the same for this tester and the sanity test, so read those sections first for tips I will not bother to repeat here.

8.1 Setting up and installing the tester

First, you need to get the tester tarfile. You can get it from the file releases on sourceforge, or, if you are using CVS, you can checkout the `AtlasTest` module. You then untar this guy in the directory you want it (`bunzip2 -c atlas_test.tar.bz2 | tar xvmf -`).

Now, you need to edit `AtlasTest/Makefile`, and set `TSTdir` to the directory you just created.

Some of these tests need a reference BLAS library to compare against, so you need to fill in your ATLAS install's `ATLAS/Make.ARCH` with a trusted, complete `BLASlib`. See the following section for details on this.

You are now ready to start the testing, as described in the following sections.

8.2 Getting a good BLASlib

Some of these tests need a reference BLAS library to compare against, so you need to fill in your ATLAS install's `ATLAS/Make.ARCH` with a trusted, complete `BLASlib`. You want an optimized library if one is available, because the Level 3 tests can go on for days extra if you use the reference library. However, only a few vendor libraries supply all of the BLAS

that ATLAS provides (to be fair, ATLAS provides BLAS above those mandated by the standard; it provides all the routines present in the Fortran77 reference library). So, the easiest way to get a complete library is to also install the reference Fortran77 library from netlib.

Get the BLAS reference tarfile from www.netlib.org/blas/blas.tgz. then do something similar to the following:

```
mkdir FBLAS
cd FBLAS
gunzip -c ../blas.tgz | tar xvf -
f77 -O -c *.f
ar r libfblas.a *.o
```

You may need to substitute for your Fortran77 compiler and flags, and if your system uses ranlib, run that on `libfblas.a` as well.

Now, you can set `BLASlib` so that the optimized library is linked in first, and the reference BLAS are used for any routines not provided in the optimized library. For instance, here's my `BLASlib` for using MKL:

```
BLASlib = /home/rwhaley/lib/libmkl32_def.a /home/rwhaley/lib/libfblas.a
```

For many routines, the tester cannot tell the difference between an error in the BLAS given by `BLASlib`, and an error in ATLAS. Subsequent section will explain how to figure this out, but understand that a lot of optimized BLAS will fail this tester, in which case you need to link against the F77 BLAS instead of the optimized version of that routine. Let us say you find out that there are errors in the optimized DTRSM. In this case, you can simply link in the F77 reference DTRSM object file first to override the one in the optimized lib. So, your `BLASlib` line would then look something like:

```
BLASlib = /home/rwhaley/FBLAS/dtrsm.o \
/home/rwhaley/lib/libmkl32_def.a /home/rwhaley/lib/libfblas.a
```

Obviously, if you have more than a few routines like this, just testing against the f77 reference BLAS and taking the extra runtime is the way to go.

8.3 Running the tester

The first thing to be aware of in running the tester is that the log files it creates can take up a *lot* of space. You can kill the log files as soon as the tester finishes, but you need enough space for it to complete. The command to run the tester is fairly simple:

```
make TOPdir=<ATLAS install directory> ARCH=<arch>
```

So, an example of this would be:

```
make TOPdir=/home/rwhaley/atlas3.3/ATLAS3.3.11.3 ARCH=Linux_PPRO_2
```

As previously mentioned, however, this tester can run as long as several days. So, if you are connected to the machine with an unreliable or short- term connection, you will need to ensure it can continue to run even if you are disconnected. Under most unixes, you can do this by using the `nohup` command. For example:

```
nohup make TOPdir=/home/rwhaley/atlas3.3/ATLAS3.3.11.3 ARCH=Linux_PPRO_2 |& \
tee PPRO.out &
```

is what I use with the tcsh shell. Bourne shell users will need a different redirect command.

8.4 Finding errors

Some errors you may see on standard out, or in the log file. If you haven't seen any there, you need to scope the stored up output created by the tester. The tester puts such output files in `AtlasTest/res/ARCH`. There's a small shell script `AtlasTest/scope.sh` which takes `ARCH` as an argument, and will grep the relevant files for errors. If it finds an error, you then edit the file in question (`scope.sh` prints the file it is grepping), and find the test run that caused the error. This can take a bit what with the volume of output, but is doable if you stick with it.

Once you have the error, you need to repeat it. You can try running the exact case, but sometimes that won't do it (for instance, you have a memory error that requires you to run many cases); you then need to find a small run that does demonstrate the error.

You should then apply the normal tricks (linking to F77 BLAS instead of sys blas, having the tester call the f77 blas twice, etc) to ensure the error really is in ATLAS, before tracking the error to its source.

8.5 Tracking down errors in the bin/ testers

There are two types of bin/ testers: lapack and blas. The BLAS testers have executable names of the form

```
x<pre>l<level>blastst
```

The BLAS testers test ATLAS against a known-good implementation, so the first thing to do is make sure the error is in ATLAS, and not the known-good implementation. To do this, compile the reference BLAS from netlib, as discussed in Section 8.2, and then relink and rerun the test in question. If the error goes away, you have found an error in your known-good library, not ATLAS. If it stays, you have found an error in ATLAS, and you should track it down or report it.

See Section 7.5 for information on tracking problems in the LAPACK testers.

9 Finding a good NB for GEMM

One of the things I do most frequently with user-submitted kernels is reduce the blocking factor that the user has chosen. I often choose smaller NB than the best for asymptotic GEMM performance, and even more often choose one that does not yield the best performance in the kernel timer. To understand why, you must understand the following points, explained in turn below:

1. Better kernel timing (eg. `make ummcase` in your `<OBJdir>/tune/blas/gemm/` directory) does not always yield better total GEMM performance
2. Large NB means significantly more time in cleanup code

3. Large NB means significantly more time in unblocked application code

9.1 Better kernel timing does not always yield faster GEMM

The kernel timer (invoked by one of the `make mmcase` variants available in `<OBJdir>/tune/blas/gemm/`) tries to mimic the way ATLAS calls the kernel. However, it does not do everything the same way. First, there is no cleanup, so it is always calling the kernel only. More importantly, `CacheEdge` has not yet been determined, so no Level 2 Cache blocking is being performed. Therefore, it may sometimes look like you are better off to block the kernel for the L2 when using these kernel timers, when in fact, if you instead block for the Level 1 cache, `CacheEdge` will then further speed things up later, and thus the smaller NB achieves better GEMM performance, even when it runs slower in the kernel timer.

For machines with very large L1 caches, often several blocking factors that fit into L1 have roughly the same performance. In such a case, it is very likely that you want to choose the smallest achieving that rough performance, as it will allow more blocks to fit into the L2 blocking to be done later.

If a kernel appears to get much better performance with a large NB, the best idea is to build a full GEMM using both the best-performing small NB, and the best performing large NB, and seeing what the gap truly is. Very often, the small kernel will actually be better even asymptotically, and if it is not, it will often be so much better for smaller problems that it makes sense to use it anyway.

Even beyond these explanations, it is sometimes the case that the kernel timer predicts good performance that is not realized when the full GEMM is built. This is usually due to inadequate cache flushing, leading to overprediction of performance because things are retained more in the cache than they are in practice. Therefore, I usually pump up the flushing mechanism (set `L2SIZE` of your `Make.ARCH` to ridiculously large levels). No matter what, actual full GEMM performance is the final arbiter. If it is not as high as predicted by the kernel timer, it may be worthwhile to see if other, smaller NB, cases achieve the same full-gemm performance.

9.2 Large NB means more time in cleanup

One bad news about choosing a large NB is that applications will spend more of their time in cleanup. Let us say you choose a block factor of 120. In this case, many applications will never even call your optimized kernel, but spend all their time in GEMM cleanup. Some applications are statically blocked, and if their NB is smaller than yours, they can spend their entire time in cleanup even for large problems.

Therefore, if you must choose a large NB in order to get adequate GEMM performance, you must pay an unusual amount of attention to cleanup optimization. However, as the next section will discuss, even if cleanup ran at the same speed as your best kernel, this will yield poor performance for many codes.

9.3 Large NB means more time in unblocked application code

Probably the worst thing about choosing a large NB is that many applications use Level 1 and 2 BLAS in order to do the unblocked part of the computation. These BLAS are usually at least an order of magnitude slower than GEMM. Therefore, as you increase NB,

for applications with unblocked portions, you increase the proportion of time spent in this order-of-magnitude slower code. Therefore, even with perfect cleanup, a large NB may result in an application running **at less than half speed**, even though GEMM performance is quite good.

To get an idea of this, simply scope the factorizations provided by LAPACK. These applications are statically blocked, so that the column factorizations (eg., DGETF2 for LU) are used until NB is reached. If ILAENV returns a blocking factor smaller than your GEMM, the applications will stay in cleanup even for large problems. Even worse, some applications (eg., QR) require workspace proportional to NB, and since dynamic memory is not used, it is possible even if you hack ILAENV to use the correct blocking factor, they will be forced to a smaller one.

9.4 Finding a good NB

I will call the first level of cache accessed by the floating point unit the Level 1 cache, regardless of whether it is the first level of cache of the machine (there are a number of machines, such as the P4 (prescott) and Itanium where the FPU skips the Level 1 cache). Let N_e be the number of elements of the data type of interest in this cache. If this cache is write-through, then a rough guess for a good upper bound is $N_B \leq \sqrt{N_e}$. If the cache is not write-through, this is still the upper bound, but many larger caches often benefit from using a smaller N_B , one roughly $N_B < \frac{\sqrt{N_e}}{2}$. We can describe this more exactly, but these bounds are easy to compute during tuning.

You should not choose an N_B that is a power of 2, as this can occasionally cause disastrous cache conflicts. There's often a small advantage to choosing N_B that are a multiple of cache line size; this can sometimes be critical, depending on the arch.

So, the basic idea is to start looking at N_B given by the above two computations, and then try a little smaller and larger using the kernel timer. If you get two that tie for out-of-cache performance, always take the smaller. If best performance is achieved with very large N_B (say $N_B \geq 80$), then always confirm that it yields better GEMM performance than a smaller N_B , and that application performance is not severely impacted, particularly for smaller problems.

The way I usually time application performance is to time ATLAS's LU. This actually gives you a very rosy picture of how a large block factor will effect performance, in that it uses recursion rather than statically blocking. This means that ATLAS's LU does not have any unblocked code, and thus doesn't slow down the way LAPACK's LU will for large N_B . However, if even this code shows performance loss for smaller sizes, you know your cleanup needs to get a *lot* better, or you need to reduce N_B , even if it results in a slight reduction in GEMM performance. If you want to get a better idea of how most applications will perform, time one of LAPACK's factorizations instead.

Under no circumstances should you choose a blocking factor much larger than 120. I confine the ATLAS search to a maximal size of 80 for the above reasons, but occasionally go a little higher for machines without effective L1 caches. However, this can absolutely kill application performance. Further, it is never a good idea to completely fill an Level 2 cache with your block. It may look good in GEMM, but it will die in any application, both for the reasons above, and the following: The L2 cache is shared instruction/data. Filling it with data will often lead to instruction loading/flushing cycle when a larger application is

calling. Remember that GEMM is of interest because of all the applications that are built from it, not when used in isolation.

If a NB larger than 60 only gives you a few percent, always choose a smaller one; only go above 80 for significant advantage, and essentially don't go above 120 unless absolutely necessary, and then you can expect slowdown in many applications, even once you have fully optimized all cleanup cases.

10 Information on atlconf

For ATLAS 3.7.12, ATLAS's configure routine was completely rewritten for greater modularity. The total amount of code probably increased, but the amount that must be examined at any time should be very much smaller.

In the new system, the topmost unit is `ATLAS/configure` which is a BFI shell script which allows ATLAS's `config.c` to be invoked in a way very similar to `gnu configure`. This shell script gathers some info and fills in a `Makfile` which is then used to build `xconfig` from `ATLAS/CONFIG/src/config.c`. `config.c` is a driver program that first calls various probes to determine any information not overridden by user flags, and then calls `xspew` to create a full `Make.inc` for the target architecture. `xspew` is built from the file `ATLAS/CONFIG/src/SpewMakeInc.c`.

The idea is to change ATLAS's install so it consists of the following commands:

1. `/path/to/ATLAS/configure` : Create `Make.inc` and build subdirs in the present directory (ATLAS no longer requires building in arch-spec directories under the source tree)
2. `make build` : Build ATLAS
3. `make test` : run sanity tests
4. `make time` : run simple benchmarks, compare observed vs. expected performance, and issue warning if too low
 - Keep record of arch default installs with performance as % of clock rate (not peak!)
 - If arch def used wt bad compiler, this will detect performance difference, and user should be warned
 - If no arch defaults (so expected performance unknown), still create standard benchmark output file for submission to ATLAS
5. `make install` : copy libraries and include files to user-specified directories
 - Will need to figure out what-all include files to copy!

10.1 Weaknesses in configure

1. Need flag `--help`:
 - Echos configures own flags, then says "here's the atlconf flags", and invokes `xconfig` with `--help`
2. Need to add `gnu configure` lookalike for installing the libraries and include files (not sure what-all to put here, maybe `--prefix`, and probably others – look up in `gcc`).

10.2 Weaknesses in spew/config

1. Need to figure out correct handling when using `ATLwin_` compiler wrapper
2. Needs a flag for the `delay` variable
3. Needs correct lib setup:
 - Spew & config flags for BLASlib, FBLASlib, FLAPACKlib, extra syslibs.
 - config.c needs an updated GetSysLib, output passed to spew via extra syslibs.
 - config needs updated GetPmake and FindTarZip, and then both routs need flags to indicate these things
 - Need flags for archiver, archflags, and ranlib (also table lookup in config.c)

10.3 Probe Overview

From ATLAS 3.7.12 on, ATLAS's config routine was rewritten for greater modularity, with each config probe having its own driver and so on. For this discussion, we will refer to the machine doing the cross-compilation as the *frontend* (abbreviated as FE), and the machine which ATLAS is being tuned for the *backend* (abbreviated as BE). Note that if you are not doing cross-compilation (the majority of the time) the front-end and back-end are the same machine.

Every type of probe has a frontend driver (occasionally, config may directly call the backend driver, if there is only one) which will itself call multiple backend drivers. For instance, the probe to compute the architecture runs on the frontend, and calls different backend drivers depending on the assembly dialect and operating system of the backend. The files for the frontend drivers are located in `ATLAS/CONFIG/src`, and the backend files are in `ATLAS/CONFIG/src/backend`, with all include files in `ATLAS/CONFIG/include`. All frontend probes use the file `atlconf_misc.c` (prototyped in `atlconf_misc.h`), which handles things like file I/O, issuing shell commands, etc. The current probes used by config are:

1. OS Probe

Purpose: Discover the Operating System being used

Inputs: None

Outputs: Enumerated type of OS

FE files: `probe_OS.c`

BE files: None (uname on BE)

2. Assembly dialect probe

Purpose: Discover what ATLAS assembly dialect works

Inputs: OS enum (gives subdialect of assembler)

Outputs: Enum of assembly dialect

FE files: `probe_asm.c`,

BE files: `probe_this_asm.c` – [`probe_gas_parisc.S`, `probe_gas_ppc.S`, `probe_gas_sparc.S`, `probe_gas_x8632.S`, `probe_gas_x8664.S`]

3. Vector ISA extension probe – assembly

Purpose: Discover which of supported vector ISA extensions work

Inputs: enums for OS and assembly dialect

Outputs: `iflag = ((1<<ISA0) | (1<<ISA1) | ... | (1<<ISAn))`

FE files: `probe_vec.c`

BE files: `probe_svec.c` – [`probe_AltiVec.S`, `probe_SSE.S`], `probe_dvec.c` – [`probe_SSE2.S`], `probe_dSSE3.c` – [`probe_SSE3.S`]

4. **Vector ISA extension probe – C:** Write this later, using C-inline statements for platforms where we don't speak the assembly, but can still use peter's vector include file

5. Architecture probe

Purpose: Discover target architecture/machine

Inputs: OS and assembly enums [force 64/32 bit usage]

Outputs: enum of arch

FE files: `archinfo.c`

BE files: `archinfo_x86.c`, `archinfo_linux.c`, `archinfo_freebsd.c`

Notes: See Section 10.4 for more details.

6. 64-bit probe

Purpose: Discover if arch supports 64-bit pointers

Inputs: OS, arch [user choice]

Outputs: 32 / 64

files: Config directly calls `archinfo`

Notes: : New policy: config assumes whatever compiler gives you w/o `-m32 -m64`, and user must throw special flag to append these to the line.

7. Compiler probe

Purpose: Find good compilers

Inputs: OS, arch [,suggested compilers]

Outputs: The following:

- (a) `F2CNAME`, `F2CINT`, `F2CSTRING` enums
- (b) Compilers and flags

FE files : `probe_comp.c` `probe_f2c.c` `probe_ccomp.c`

BE files : `f2cname[F,C].[f,c]`, `f2cint[F,C].[f,c]`, `f2cstr[F,C].[f,c]`, `ccomp`
interaction not yet done

Note: This is complex, see Section 10.6 for details.

8. Arch defaults probe

Purpose: Discover arch defaults

Inputs: OS, arch, compilers

Outputs: Whether to use arch defs (INSTFLAGS in `Make.inc`

files: ARCHS/Makefile

invoke: Arch default setup is instigated by `atlas_install.c`.

notes: May want to have it autobenchmark kernel, test against table of expected perf, to see if arch def are OK wt this compiler version.

10.4 Architectural Probes

We use the `archinfo_xxx` probes to discover the following architectural information:

- 'n': number of cpus
- 'c': number of cache levels
- 'C' #: size in KB of cache level #
- 'v': verbose (prints strings as well as ints)
- 'm': clock rate in Mhz
- 'a': ATLAS architecture classification
- 'b': support for 64 / 32 bits
- 't': Is cpu throttling currently on (-1: no, 0: don't know, 1: yes)

If a given probe cannot find that particular item, it is returned as 0.

The frontend wrapper script `archinfo.c` calls these guys according to OS, and tries to get all flags filled in with union of functionality of `archinfo_x86` and `archinfo_<OS>`.

The following probes run on the backend. We have the x86-specific probe `archinfo_x86.c`, which is designed to work on any x86-using OS (though assembly problems could prevent it working on unknown architectures). All other probes query the OS for information, and these are:

- `archinfo_linux.c`: Basics work!
- `archinfo_freebsd.c`: Basics work.
- `archinfo_solaris.c`: *test on USIII/USIV – not yet implemented*
- Eventually should add irix and hpux; I think still have irix access from todd, and hpux may still do testdrive?
 - Irix & hpux now primarily Itanium, so see if we can get `archinfo_x86` working under linux/itanium before considering these platforms

10.5 Notes on configure

New policies:

- Any 64-bit arch defaults to building whatever the compiler natively does w/o flags unless overridden
- L2SIZE always set to 4MB unless overridden
- Timer defaults to standard wall/cpu unless overridden
 - Exception is solaris, where we default to solaris hr timers
 - If Mhz is passed in, use cycle-accurate wall-timer on x86
- Need to probe for CPU throttling and abort if true!
- For deprecated machines/OSes I don't have access to, ask dev list to post config patch if they want support

Deprecated machines (no longer supported in config or arch def):

- MIPS chips (add back in using todd's machines – common embedded?)
- Alpha chips (some EV6 may still be alive – UTK?)
- All pre-G4 PowerPC chips
- All pre-UltraSPARC Sun chips

Going to start out with Linux and FreeBSD(OS X) support. Add them in following order:

1. Solaris
2. Windows
3. Irix
4. HP/UX

10.6 Compiler Setup and Handling in ATLAS Config

This is complicated as hell. Potentially, each architecture/OS combo has unique compiler and flags for each supported compiler (more below), and the user can override any/all of these. I'm changing the number of supported compilers for greater flexibility. These are:

ICC : compiles all interface/non-computation routines, and particularly any routine that does I/O. This should be the compiler that the user compiles his own code with, but it may not be good enough to compile ATLAS's kernels with. This compiler's performance could have some slight impact on performance, particularly if it doesn't do loops well, but since it is not used for any kernel compilation the impact should be minimal.

SMC : used to compile ATLAS single precision matmul kernels

DMC : used to compile ATLAS double precision matmul kernels

SKC : used to compile ATLAS single precision non-matmul computation kernels (eg., l1 and l2 kernels)

DKC : used to compile ATLAS double precision non-matmul computation kernels (eg., l1 and l2 kernels)

XCC : used to compile all front-end codes

F77 : Valid fixed-format Fortran77 compiler that compiles ATLAS's F77 interface routines. This should match the Fortran77 the user is using. This compiler's performance does not affect ATLAS's performance, and so may be anything.

Here's my present design:

1. **Compiler defaults:** are read in from `atlcomp.txt`, which allows the user specify default compiler/flags, as well as specific ones for particular architectures, and multiple compilers for a given arch.
 - `ATLAS/SRC/probe_comp.c`
2. **Executable search:** takes name of executable (in this case a compiler name), and finds the path to it. Skipped if the user provides the path as part of the compiler.
 - *unimplemented:* presently use compilers w/o path
3. **C compiler interaction probe:** separate probe that takes two or more C compilers and their flags as arguments, and makes sure they are able to call each other w/o problems.
 - *unimplemented:* presently assumed to work
4. **F77/C calling convention probe:** as in present config, but as an independent probe.
 - *done:* front-end is `probe_f2c.c`

Some notes on compiler handling:

- Config will die if no f77/C compiler found, with message to user to pass in correct compiler.
- Have flag indicating build atlas w/o a fortran compiler.
- Have flag for some common cases:
 - MSVC++ for icc, but gcc for rest
 - gcc 4 icc, gcc 3 rest
 - append `-m32/64` to all flags

- Use icc for interface or everything (whichever faster)
- Use xlc for interface or everything (POWER support, I guess?)
- Have ability to append user flags to atlas's default flags (works in addition to -m32/64 appending)
- ???