


# Architectures vs the Ports tree: a losing battle ?

Marc Espie <espie@openbsd.org>, <espie@lse.epita.fr>


OpenBSD project & Laboratoire de Recherche de l'Epita

September 18, 2022



# Elephant





## Only on OpenBSD

- we only use cross-compilation for bootstrap
- every arch builds its own packages
- best stress-test ever

- strict alignment architectures
- big endian vs little endian
- character signedness (not really interesting)
- reverse stack
- ghostguard
- smallkva
- (compiler bugs)

- Cumulative work over the past 20 years or so
- Lots of (smallish) topics I haven't talked about ever
- No big plan, just lots of small improvements and know-how

# What's an architecture

## What's in a name

- ARCH describes the exact machine (e.g., macppc)
- MACHINE\_ARCH is the "cpu make" (e.g., powerpc)
- details like "i386" vs "pentium" are generally not encoded
- → generally, packages target MACHINE\_ARCH

## Compiler subversion

- Compilers offer `-march=native` options
- This should **never** be used for building packages
- Instead the base OS targets a baseline cpu, and everything should work on this cpu and later versions
- (notable exception: the altivec extensions to ppc, a while ago)
- slowly, the bar gets raised, from i386 to i586 to...



## talking to upstream

- explaining that we're software vendors, and we need reliable builds that will work on every machine
- so no tests during builds to optimize the compilation to the exact machine we have
- specifically for multimedia software: no hand-crafted assembly code selected at runtime

## good practices

- provide at least a way to build that doesn't hardcode machine details
- replace compile-time tests with runtime tests to select hand-crafted code (for instance, relying on cpuid on intel boxess)

## Compiler options

- in general, upstream is bad with compiler options
- those do break on some arches
- so we standardize on `-O2` and `-O2 -g`
- porters try to help heeding `CFLAGS` and `CXXFLAGS`
- we hate build systems without an easy way to specify options
- even compilers change options with hilarious effects



- coding tests on `MACHINE_ARCH` is an extraordinarily bad idea
- prefer `ONLY_FOR_ARCHS` and `NOT_FOR_ARCHS`
- (or eventually `BROKEN`)
- that way everything is referenced properly

## dpb

- gets information through `make dump-vars`
- should be resilient to errors
- will flag as errors missing information for ports
- removes stuff if marked as not available for this arch
- can even be run on a different architecture for listing

## sqlports

- we also run `make dump-vars` to create a db of everything
- that one errors out if something does not work, possibly a `pkgpath`

# example 1

```
1  ==> archivers
2  ==> archivers/arc
3  archivers/arc.IS_INTERACTIVE=No
4  archivers/arc.SUBPACKAGE=-
5  archivers/arc.BUILD_PACKAGES= -
6  archivers/arc.MULTI_PACKAGES=-
7  archivers/arc.DISTFILES=arc-5.21p.tar.gz
8  archivers/arc.MASTER_SITES=https://downloads.sourceforge.net/sourceforge/arc/
9  archivers/arc.CHECKSUM_FILE=/usr/ports/archivers/arc/distinfo
10 archivers/arc.FETCH_MANUALLY=No
11 archivers/arc.PERMIT_DISTFILES=Yes
12 archivers/arc.NO_TEST=Yes
13 archivers/arc.TEST_IS_INTERACTIVE=No
14 archivers/arc.DISTNAME=arc-5.21p
15 archivers/arc.HOMEPAGE=http://arc.sourceforge.net/
16 archivers/arc.MAINTAINER=The OpenBSD ports mailing-list <ports@openbsd.org>
17 archivers/arc.USE_GMAKE=No
```



## example 11

```
18 archivers/arc.USE_GROFF=No
19 archivers/arc.NO_BUILD=No
20 archivers/arc.USE_LIBTOOL=Yes
21 archivers/arc.SEPARATE_BUILD=No
22 archivers/arc.TARGETS= do-install
23 archivers/arc.MAKEFILE_LIST=/usr/share/mk/sys.mk Makefile /usr/share/mk/bsd.port.mk
24 archivers/arc.USE_LLDB=Yes
25 archivers/arc.USE_WXNEEDED=No
26 archivers/arc.COMPILER=base-clang base-gcc gcc3
27 archivers/arc.COMPILER_LANGS=c c++
28 archivers/arc.COMPILER_LINKS= clang /usr/bin/clang clang++ /usr/bin/clang++ cc /usr/bin/cc
29 archivers/arc.SUBST_VARS=ARCH BASE_PKGPATH FLAVOR_EXT FULLPKGNAME HOMEPAGE LOCALBASE
30 archivers/arc.PKGPATHS=archivers/arc
31 archivers/arc.FULLPKGNAME=arc-5.21pp0
32 archivers/arc.PERMIT_PACKAGE=Yes
33 archivers/arc.COMMENT=create & extract files from DOS .ARC files
34 archivers/arc.PKGNAME=arc-5.21p
```



## example III

```
35 archivers/arc.PKGSPEC=arc-*
36 archivers/arc.PKGSTEM=arc
37 archivers/arc.PREFIX=/usr/local
38 archivers/arc.WANTLIB=c
39 archivers/arc.CATEGORIES=archivers
40 archivers/arc.DESCR=/usr/ports/archivers/arc/pkg/DESCR
41 archivers/arc.REVISION=0
42 archivers/arc.STATIC_PLIST=Yes
43 archivers/arc.PKG_ARCH=amd64
44 ==> archivers/blosc
45 archivers/blosc.BUILD_DEPENDS=devel/cmake devel/ninja
46 archivers/blosc.IS_INTERACTIVE=No
47 archivers/blosc.SUBPACKAGE=-
48 archivers/blosc.BUILD_PACKAGES= -
49 archivers/blosc.MULTI_PACKAGES=-
50 archivers/blosc.DISTFILES=c-blosc-1.21.1.tar.gz
```

- every location in the ports tree has a unique fullpkgpath
- for instance, archivers/arc or lang/python/3.10,-tests
- there are FLAVORS and MULTI\_PACKAGES

- variations are often specific parts that do not build on an architecture
- we can setup a MULTI\_PACKAGES port with that part in a separate SUBPACKAGE
- tests won't work because those subpackages won't be reachable
- so instead we remove stuff: MULTI\_PACKAGES → BUILD\_PACKAGES

# Example 1

```
1 ONLY_FOR_ARCHS-java =      aarch64 amd64 i386
2
3 CATEGORIES =               graphics devel
4 COMMENT-main =            library for computer vision real-time processing
5 COMMENT-java =            Java bindings for OpenCV
6
7 V =                        4.6.0
8 GH_ACCOUNT =               opencv
9 GH_PROJECT =               opencv
10 GH_TAGNAME =               ${V}
11
12 PKGNAME-main =             opencv-${V}
13 PKGNAME-java =             opencv-java-${V}
14
15 HOMEPAGE =                 https://www.opencv.org/
16
17 MAINTAINER =               Rafael Sadowski <rsadowski@openbsd.org>
```





## Example II

```
18
19 .for i in opencv_calib3d opencv_core opencv_features2d \
20     opencv_flann opencv_highgui opencv_imgproc opencv_ml opencv_objdetect \
21     opencv_photo opencv_stitching opencv_video opencv_imgcodecs \
22     opencv_videoio opencv_dnn
23 SHARED_LIBS += $i 10.0
24 .endfor
25
26 WANTLIB-main += ${COMPILER_LIBCXX} avcodec avformat avutil OpenEXR-3_1
27 WANTLIB-main += c cairo gdk-3 gdk_pixbuf-2.0 glib-2.0 gobject-2.0 gstapp-1.0
28 WANTLIB-main += gstbase-1.0 gstaudio-1.0 gstpbutils-1.0 gstreamer-1.0
29 WANTLIB-main += gstriff-1.0 gstvideo-1.0 gtk-3 jpeg m openjp2 png swscale tiff
30 WANTLIB-main += webp z
31
32 WANTLIB-java += ${COMPILER_LIBCXX} opencv_calib3d opencv_core opencv_dnn
33 WANTLIB-java += opencv_features2d opencv_flann opencv_imgcodecs
34 WANTLIB-java += opencv_imgproc opencv_ml opencv_objdetect opencv_photo
```



## Example III

```
35 WANTLIB-java += opencv_video opencv_videoio
36
37 COMPILER =          base-clang ports-gcc
38
39 MULTI_PACKAGES =    -main -java
40 PSEUDO_FLAVORS =    no_java
41 FLAVOR ?=
42
43 # BSDL
44 PERMIT_PACKAGE =    Yes
45
46 MODULES =           devel/cmake \
47                     lang/python
48
49 BUILD_DEPENDS =     math/eigen3 \
50                     math/py-numpy${MODPY_FLAVOR}
51
```



# Example IV

```
52 RUN_DEPENDS-main =      math/py-numpy${MODPY_FLAVOR}
53
54 RUN_DEPENDS-java =      ${MODJAVA_RUN_DEPENDS}
55
56 LIB_DEPENDS-main =      ${LIB_DEPENDS} \
57 graphics/ffmpeg \
58 graphics/jpeg \
59 graphics/libwebp \
60 graphics/openexr \
61 graphics/openjp2 \
62 graphics/png \
63 graphics/tiff \
64 multimedia/gstreamer1/core \
65 multimedia/gstreamer1/plugins-base \
66 x11/gtk+3
67
68 LIB_DEPENDS-java =      ${BUILD_PKGPATH},-main=${V}
```

# Example V

```
69
70 # XXX PIE cannot be produced due to problems with inline assembly.
71 # Since OpenCV is mostly used as a LIBrary, switch to PIC.
72 .if ${MACHINE_ARCH:Mi386}
73 CFLAGS +=          -fPIC
74 CXXFLAGS +=        -fPIC
75 .endif
76
77 CONFIGURE_ARGS =   -DBUILD_DOCS=OFF \
78                   -DBUILD_EXAMPLES=OFF \
79                   -DBUILD_IPP_IW=OFF \
80                   -DBUILD_ITT=OFF \
81                   -DBUILD_PERF_TESTS=OFF \
82                   -DBUILD_TESTS=OFF \
83                   -DBUILD_opencv_python2=OFF \
84                   -DINSTALL_PYTHON_EXAMPLES=OFF \
85                   -DINSTALL_TESTS=OFF \
```



# Example VI

```
86 -DOPENCV_SKIP_PYTHON_WARNING=ON \  
87 -DPYTHON_DEFAULT_EXECUTABLE=${MODPY_BIN} \  
88 -DWITH_1394=OFF \  
89 -DWITH_ADE=OFF \  
90 -DWITH_CUDA=OFF \  
91 -DWITH_EIGEN=OFF \  
92 -DWITH_IPP=OFF \  
93 -DWITH_OPENCL=OFF \  
94 -DWITH_V4L=ON \  
95 -DWITH_VTK=OFF \  
96 -DOPENCV_GENERATE_PKGCONFIG=ON  
97  
98 .include <bsd.port.arch.mk>  
99  
100 .if ${BUILD_PACKAGES:M-java}  
101 MODULES +=          java  
102 MODJAVA_VER =       1.8+
```



## Example VII

```
103 BUILD_DEPENDS +=          devel/apache-ant
104 .else
105 # Safe: Java will be detected, if present, but won't be used
106 CONFIGURE_ARGS +=        -DBUILD_opencv_java=OFF
107 .endif
108
109 CONFIG_ADJ_CMD =          perl -pi
110 .for _l _v in ${SHARED_LIBS}
111 CONFIG_ADJ_CMD +=        -e 's,lib${_l}.so([\^\.]),lib${_l}.so.${_v}$$1,g;'
112 .endfor
113
114 NO_TEST =                 Yes
115 # Enable to run the regression tests
116 #TEST_IS_INTERACTIVE =      X11
117 #
118 #CONFIGURE_ARGS +=         -DDBUILD_TESTS=ON \
119 #                           -DBUILD_PERF_TESTS=ON
```



## Example VIII

```
120
121 post-patch:
122     perl -pi -e 's@^.*(#\s*include)@$$1@' \
123         ${WRKSRC}/samples/cpp/tutorial_code/core/how_to_scan_images/how_to
124
125 post-install:
126     ${MODPY_BIN} ${MODPY_LIBDIR}/compileall.py ${WRKINST}${MODPY_SITEPKG}
127
128 do-test:
129     cd ${WRKBUILD}; \
130     ${MODPY_BIN} ${WRKSRC}/modules/ts/misc/run.py
131
132 .include <bsd.port.mk>
```

- a part of `bsd.port.mk`
- if you don't include it yourself, it will be done automatically
- set up `BUILD_PACKAGES` according to `PSEUDO_FLAVORS` and arches
- then you test according to `BUILD_PACKAGES` for configure tests



## also properties |

```
1  # architecture constants
2
3  ARCH ?!= uname -m
4
5  ALL_ARCHS = aarch64 alpha amd64 arm arm64 armv7 hppa i386 landisk loongson \
6             luna88k m88k macppc mips64 mips64el octeon powerpc64 riscv64 sgi \
7             sh sparc64
8  # normally only list MACHINE_ARCH (uname -p) names in these variables,
9  # but not all powerpc have apm(4), hence the use of macppc
10 APM_ARCHS = arm64 amd64 i386 loongson macppc sparc64
11 BE_ARCHS  = hppa m88k mips64 powerpc powerpc64 sparc64
12 LE_ARCHS  = aarch64 alpha amd64 arm i386 mips64el riscv64 sh
13 LP64_ARCHS = aarch64 alpha amd64 mips64 mips64el powerpc64 riscv64 sparc64
14 GCC4_ARCHS = alpha hppa sh sparc64
15 GCC3_ARCHS = m88k
16 # XXX easier for ports that depend on mono
17 MONO_ARCHS = aarch64 amd64 i386
```

## also properties II

```
18 OCAML_NATIVE_ARCHS = aarch64 amd64 i386
19 OCAML_NATIVE_DYNLINK_ARCHS = aarch64 amd64 i386
20 GO_ARCHS = aarch64 amd64 arm armv7 i386 mips64
21 RUST_ARCHS = aarch64 amd64 i386 powerpc64 riscv64 sparc64
22
23 # arches where the base compiler is clang
24 CLANG_ARCHS = aarch64 amd64 arm i386 mips64 mips64el powerpc powerpc64 riscv64
25 # arches using LLVM's linker (ld.lld); others use binutils' ld.bfd
26 LLD_ARCHS = aarch64 amd64 arm i386 powerpc powerpc64 riscv64
27
28 # arches where ports devel/llvm builds - populates llvm ONLY_FOR_ARCHS
29 # as well as available for PROPERTIES checks.
30 LLVM_ARCHS = aarch64 amd64 arm i386 mips64 mips64el powerpc powerpc64 riscv64 sparc64
31 # arches where ports-gcc >4.9 exists. To be used again for modules
32 GCC49_ARCHS = aarch64 alpha amd64 arm hppa i386 mips64 mips64el powerpc powerpc64
33
34 MODGCC4_VERSION?=8
```



```
35 # arches where there is a C++11 compiler, either clang in base or ports-gcc  
36 CXX11_ARCHS = ${CLANG_ARCHS} ${GCC49_ARCHS}  
37 DEBUGINFO_ARCHS = aarch64 amd64
```

- 9700 Makefiles and fragments
- 200 uses of `bsd.port.arch.mk`
- 90 tests on `BUILD_PACKAGES`

Stuff like this actually works:

```
1 ONLY_FOR_ARCHS-sub = ${RUST_ARCHS}
2
3 .include <bsd.port.arch.mk>
4 .if ${BUILD_PACKAGES:M-sub}
5 ...
6 .endif
```

- we had binary packages in 2000
- dpb dates back from 2010
- dedicated build farms for most architectures
- takes between 24 hours and a few weeks
- regular build stats for everything (thanks landry@)

- intel 64 bits acts as "the bellwether" (most stuff always builds)
- other architectures get fixed depending on needs
- some big stuff is (sometimes) not even built because of practicality

- this was painful to create but works
- there's a variable `COMPILER` you can set to choose "the best" compiler
- some systems have `gcc3` in base, others have `gcc 4.2` and others have `clang`
- there's also a more modern `gcc` in ports and an `llvm` port
- `COMPILER` is a list of preferred compilers: `base-gcc`, `base-clang`, `gcc3`, `ports-gcc`, `ports-clang`
- either it's there, or it's not
- links under `WRKDIR/bin` will be created



- bootstrapping stuff like go and rust is painful
- we got a mechanism for PSEUDO\_FLAVORS to help dpb and preserve bootstrap

- lazy make: variable definitions first
- then tests and targets
- but MODULES
- but COMPILER
- but `bsd.port.arch.mk`
- very specific location (best of both worlds)

- language support is the #1 problem (modern C++, rust, go)
- 32 bit arches are losing
- we got dpb annotations to help (lonesome) but it's still a problem

???

Any questions ?