Browsing and Hacking the Linux Kernel with KDevelop

Alexandre Courbot, NVIDIA

June 8, 2012

About This Guy...

- Former academic in embedded OSes
- KDE user since 2003
- Employed by NVIDIA as a Tegra CE engineer
- Need to browse, study and understand kernel code quickly...
- ...and I know I am not alone

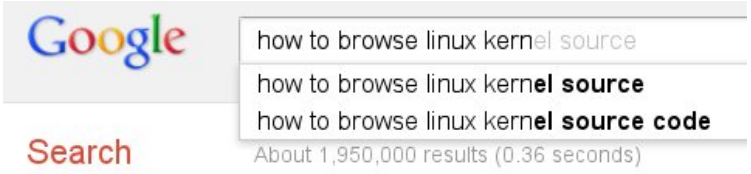

My Journey with Code Editors

Like most people, I spent some time using the most popular code editors:

- Emacs user for 5 years
- Vim user for 4 years

Both are very nice and powerful. But the '80s called, and they want their UI back.

Why are we still using tools that are between 15 and 30 years old to browse and edit kernel code?

CTags/CScope

Supported by the kernel Makefile:

```
$ make tags
$ make cscope
```

- Output usable by virtually every editor
- Do a good job at filtering arch-relevant files...
- ...but unfortunately include all mach
- ...and not `.config`-aware
- No incremental update of tags
- Takes 40-seconds on this machine when kernel source is cached
- Sometimes not so smart

LXR, Linux Cross Reference

A HTML, browseable version of the Linux source code (`http://lxr.linux.no`).

- Web-based, not integrated into your editor
- Works with some reference source, not *yours*
- Parses *everything* and returns references about everything
- Parsing results sometimes totally wrong: every `DEFINE_MUTEX` use is interpreted as a re-declaration!

CEDET (Emacs IDE)

Emacs-only Integrated Development Environment.

- Code completion
- Code browser
- ASCII UML diagrams
- ...

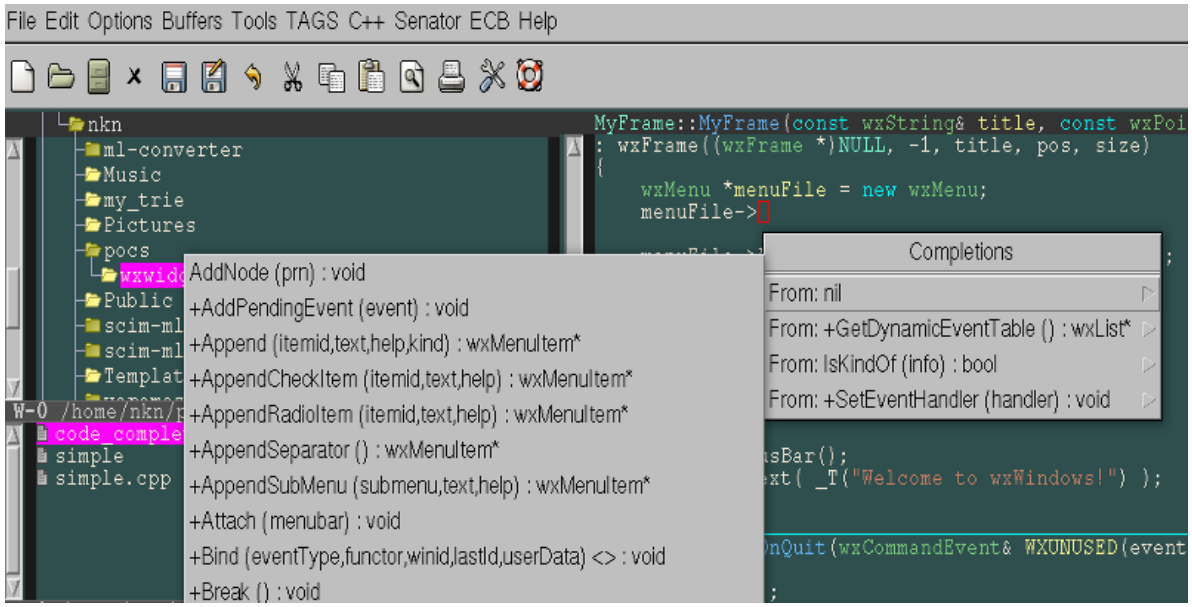

CEDET (Emacs IDE)

A "Gentle" Introduction to CEDET

```
(defun my-cedet-hook ()
  (local-set-key [(control return)] 'semantic-ia-complete-symbol)
  (local-set-key "\C-c?" 'semantic-ia-complete-symbol-menu)
  (local-set-key "\C-c>" 'semantic-complete-analyze-inline)
  (local-set-key "\C-cp" 'semantic-analyze-proto-impl-toggle))
(add-hook 'c-mode-common-hook 'my-cedet-hook)
;; automatic name completion
(defun my-c-mode-cedet-hook ()
 (local-set-key "." 'semantic-complete-self-insert)
 (local-set-key ">" 'semantic-complete-self-insert))
(add-hook 'c-mode-common-hook 'my-c-mode-cedet-hook)
```

Now open a file, and admire your screen frozen for one minute.

What To Expect From a Modern IDE?

Things have changed since last century. We have multi-core machines, filesystem notifiers, modern UIs, SSD disks.

- Take advantage of modern hardware and kernel features: `inotify`, multi-core, ...
- Background parsing
- Incremental updates of modified files
- Easy to setup, minimal configuration
- Integrated into a decent editor
- Modern UI yet allow both keyboard and mouse control

Eclipse does that well for Java - why don't we have the same for C and kernel?

About KDE

A modern, slick set of libraries, core services, and applications based on Qt.

- Highly-configurable
- Reusable components (KParts)
- DBUS-controllable
- Very dynamic and listening community
- Kate text editor
  - Syntax highlighting for 200+ files types
  - Code folding
  - Split screen
  - Very configurable, extensible through plugins and scripts
  - VI mode!

About KDevelop

Project started in 1998, went through several rewrites.
KDevelop 4:

- In development since 2005, first stable release in 2010
- Leverages most KDE technologies
- Extremely modular and extensible
- Very powerful code analysis/browsing capabilities
- Probably the most overlooked Linux/C++ Linux IDE

What happens if we open the kernel source with it?

KDevelop and the Kernel

Most features working out of the box, but:

- KDevelop's parser is C++ **only**
- Code parsing extremely long, many unnecessary files parsed
- Include paths incorrect
- No kernel configuration awareness

Tuning KDevelop for the Kernel

Two-fronts work:

1. Make KDevelop more kernel-friendly
   - Make the parser capable of handling pure C
2. `kdev-kernel` plugin
   - Kernel project manager:
     - Parse configuration and Makefile to only consider source files relevant to the current configuration
     - Declare configuration macros to guide the C parser
     - Setup the correct include paths
   - Kernel builder
     - Integrate kernel configuration GUI
     - Handle out-of-source building and cross-compiling automatically

Guided Tour

Let's see how this works!

Obvious Features

- Sessions
- Bookmarks
- Navigation history
- Customizable keyboard shortcuts
- Code snippets
- External scripts
- Customizable indentation rules
- Working files sets

Kernel Project Configuration

Kernel project configuration dialog: just choose your architecture and base configuration, and there you go!

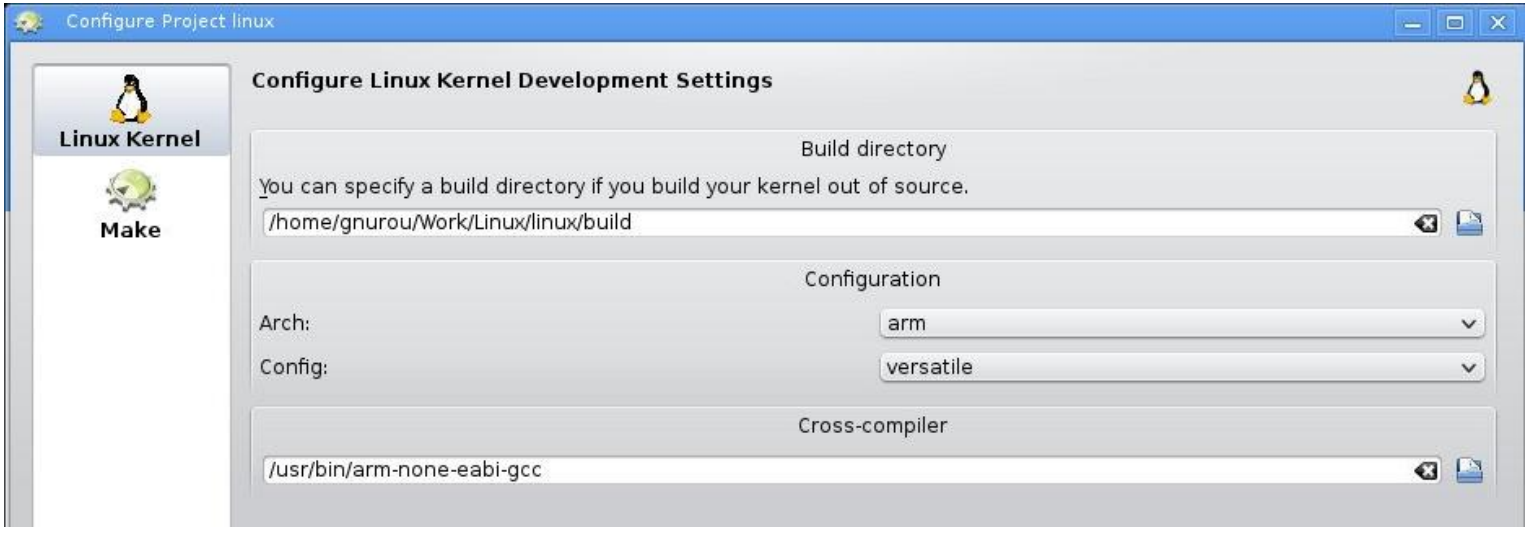

Code parsing then becomes configuration-aware

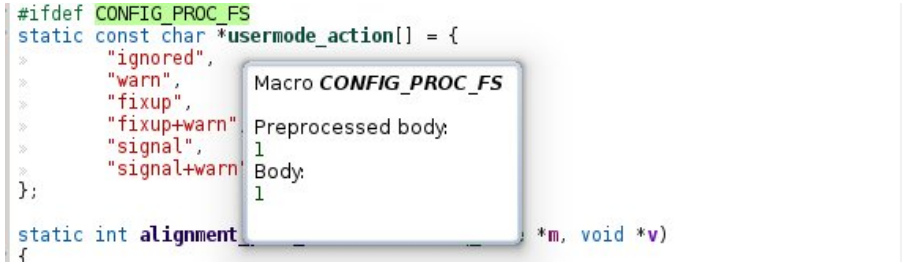

Inline Symbols Information

Links to symbols definitions, inline documentation

Macro expansion and quick peeking

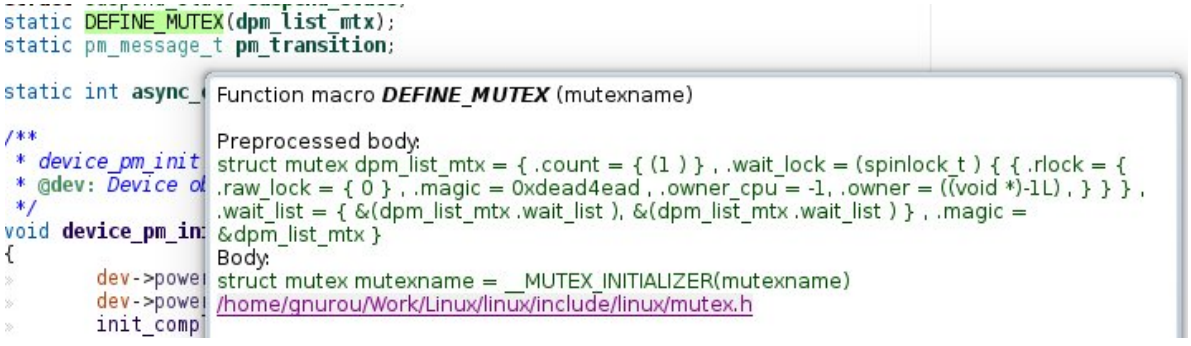

Project-wide Uses of Symbols

Find exactly where a given symbol is used

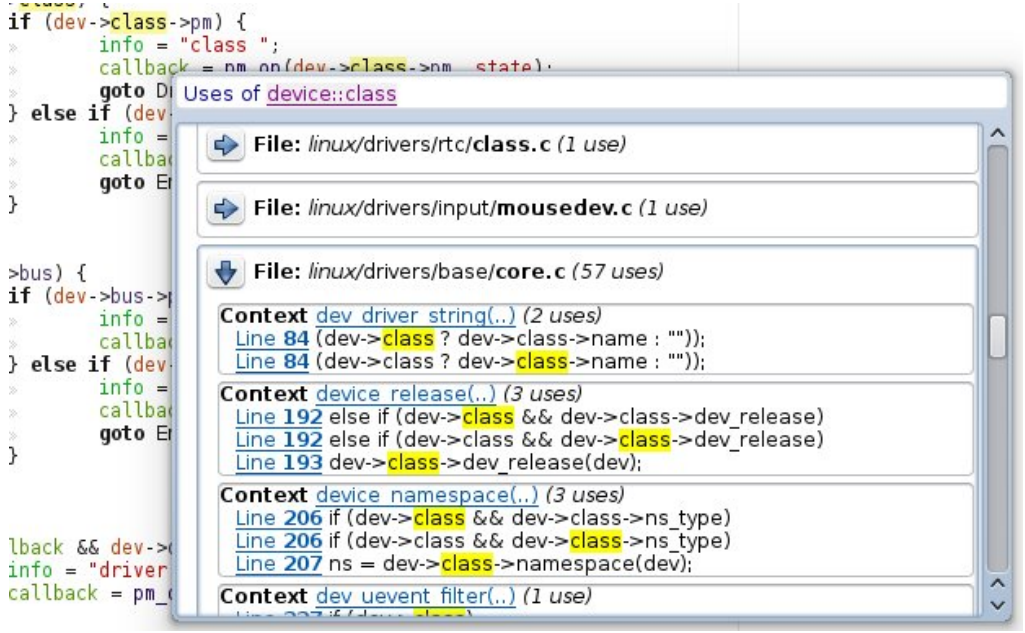

Quick Open

Quickly find any file, function of struct across the project

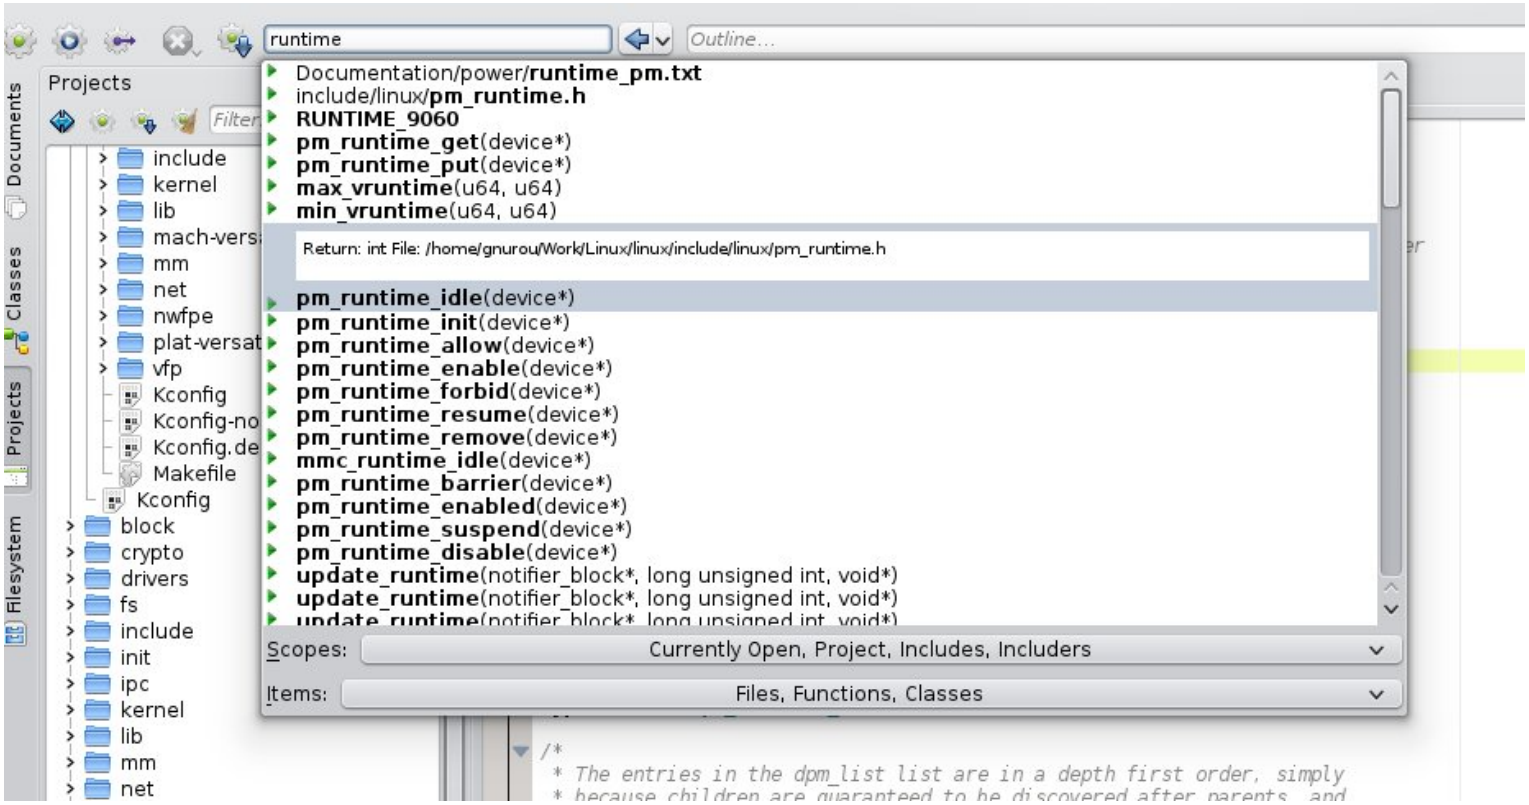

Code Editing

Smart code completion

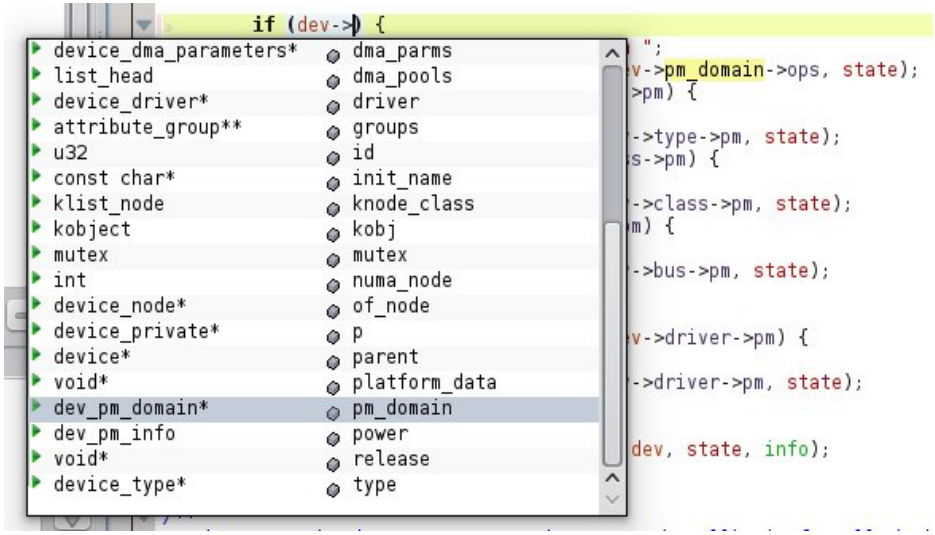

Refactoring

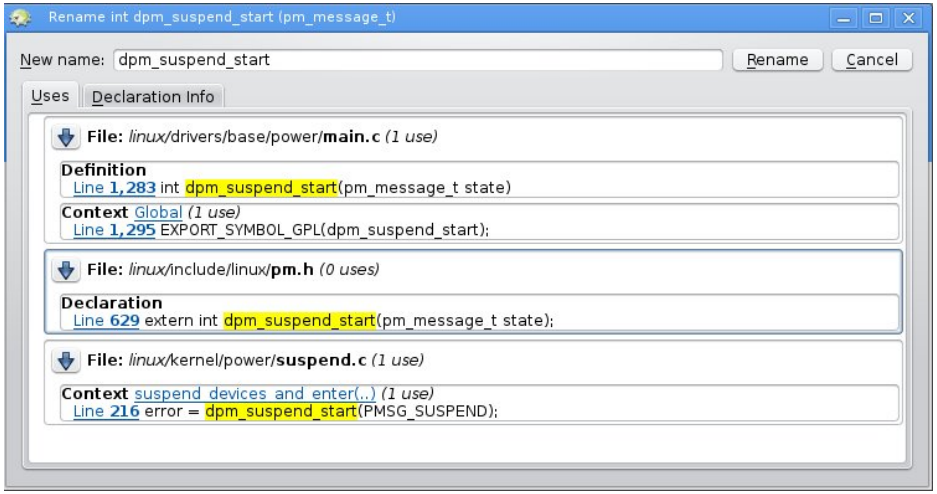

Integrated Configuration and Compilation

Integrated `make` invoked with the right parameters and linked output

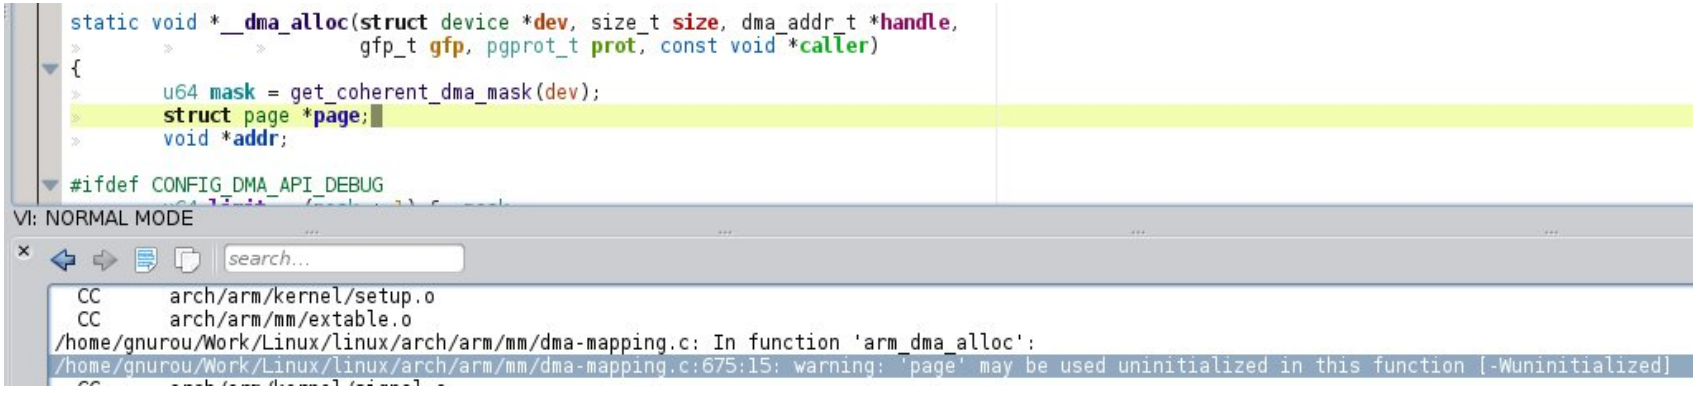

Git Integration

`git blame` support

Future Work

Both KDevelop and `kdev-kernel` are works in progress.

- Complete C support / Fix parser errors
- Improve background parsing speed
- Function pointers analysis
- Debugger integration
- Support for out-of-tree kernel modules
- Static analysis tools (call graphs, . . .)
- . . .

KDevelop is a fun project to hack on

- Usable on a daily basis and first-of-his-class on a lot of features
- Yet potential to implement many cool things!

Thank you!

KDevelop Official Website

- `http://www.kdevelop.org`

Kernel Tailored Branch and `kdev-kernel` Plugin

- `https://github.com/Gnurou/kdevelop`
- `https://github.com/Gnurou/kdev-kernel`

Feel free to contact me on Github if you have trouble setting up!
Bug reports, features requests, and patches are very welcome.