

---

# What Makes Software High-Quality?

Shlomi Fish <shlomif@shlomifish.org>

Copyright © 2008 Shlomi Fish

This work is licensed under the Creative Commons Attribution 2.5 License [<http://creativecommons.org/licenses/by/2.5/>] (or at your option a greater version of it).

Revision History		
Revision 1675	2007-04-02	shlomif
Forked the template from a previous work and working on it.		
Revision 1871	2008-05-04	shlomif
Finalised the first draft. About to receive feedback.		
Revision 1872	2008-05-04	shlomif
Fixed many spelling/grammar/etc. problems, trailing whitespace, etc.		
Revision 1882	2008-05-06	shlomif
Fixed many spelling, grammar and syntax problems (thanks in part to Omer Zak). Added some bolds, and id=""s, added the about section, reorganised the text a bit, and added the note about responsiveness and startup time.		
Revision 1884	2008-05-06	shlomif
Changed "a software" to something more idiomatic, and fixed a bad phrasing towards the beginning.		
Revision 5218	2008-05-15	shlomif
Forked the document from the first revision, and converted to DocBook-4.5.		
Revision 5329	2008-05-24	shlomif
Added the suggestion of speaking about "Why software quality is important", and the note about popularity to the introduction. Placed the Intro FCS story in the "Motivation" section. Added the "organisation" sub-section to the Intro. Added the note about the generic "weight function". Added a link to "The Stanford Checker". Corrected many "a software" instances. Added the note about how Linux and Windows looked to the Aesthetics. Added the note about low-quality code and organisational quality to the "good code" section. Added the footnote about a comprehensive comparison of Freecell solvers. Added the note about solution length and "other advantages" of fc-solve. Updated the "Thanks" section. Made many other corrections.		
Revision 5492	2008-06-09	shlomif
Fixed "manifest" to "is manifested".		
Revision 4859	2011-06-05	shlomif
Convert to Unicode single-quotes and double-quotes.		

## Table of Contents

Introduction .....	2
Motivation .....	2
Organisation of this Document .....	3
Parameters of Quality .....	3
The Program is Available for Downloading or Buying .....	4
The Version Number is Clearly Indicated .....	4
The Source is Available, Preferably under a Usable Licence .....	4
It "Just Works" .....	5
The Program has a Homepage .....	6
The Program is Easy to Compile, Deploy and Install .....	6

The Program Has Packages for Most Common Distributions .....	7
The Program Has Good, Extensive and Usable Documentation .....	7
Portability .....	7
Security .....	8
Backwards Compatibility .....	8
Good Ways to Provide Support .....	9
Speed and High Performance .....	9
Aesthetics .....	10
Conclusion .....	11
Ways to Achieve Quality .....	11
Having Modular, Well-Written Code .....	11
Automated Tests .....	12
Beta Testers .....	12
Frequency of Releases .....	12
Good Software Management .....	13
Good Social Engineering Skills .....	13
No Bad Politics .....	13
Good Communication Skills .....	13
Hype .....	14
A Good Name .....	14
Conclusion .....	15
Analysis of the Quality of the Various Freecell Solvers .....	15
Introduction .....	15
The Analysis Itself .....	15
Conclusion .....	17
About This Document .....	18
Copyrights .....	18
About the Author .....	18
Acknowledgements .....	18

## Introduction

Why is high quality in software important? Low-quality software applications will require the users or end-developers to work around their bugs and limitations, write a lot of extra functionality themselves, and as a result, duplicate a lot of effort and cause a lot of frustration and unhappiness. This is assuming they don't give up on it soon or right away, and end up looking for something else.

But what makes a popular software application "high-quality"? A high-quality program or library can be defined as one that is usable as is, that causes very few frustrations, and that requires little if any modifications or workarounds to get it up-to-speed.

I'm not talking about the application that's the most hyped, or the fastest, most featureful, or best. Sometimes you can often hear endless rants and Fear, Uncertainty and Doubt (FUD) [[http://en.wikipedia.org/wiki/Fear,\\_uncertainty\\_and\\_doubt](http://en.wikipedia.org/wiki/Fear,_uncertainty_and_doubt)] attacks about it. Often, it has competing programs that are superior in many respects.

In this essay we'll look at what quality metrics apply to software.

## Motivation

How did it occur to me to ponder this question? I am active in the fc-solve-discuss mailing list [<http://tech.groups.yahoo.com/group/fc-solve-discuss/>], which is dedicated to computerised techniques for solving the Freecell solitaire card game [<http://www.freecell.org/>], and for discussing other automated

solving and research of Solitaire games. This mailing list was started after I wrote Freecell Solver [<http://fc-solve.shlomifish.org/>] (with a lowercase “c” and an uppercase “S”), an open-source solver, which has proven to be quite popular. However, the mailing list was not restricted to discussing it exclusively, and was joined by many other Freecell solving experts, researchers and enthusiasts.

Recently, Gary Campbell wrote this message [<http://tech.groups.yahoo.com/group/fc-solve-discuss/message/741>] reading:

I think the solver discussion in the Wikipedia should mention that the FCPro solvers give quite lengthy and virtually unusable solutions. No human wants to follow the several hundred steps that usually result in order to get at what’s really required to solve a given layout. If someone wants a solution that can be understood, they want one that is under 100 steps. This is a major contribution and it should be stated. There are a lot of “**toy**” solvers and only a very few of “**industrial strength.**” One of the latter is the solver by Danny Jones. I don’t see that on your list. The results he has gotten from his solver pretty much put your solver to shame. The more you hype Freecell Solver, the more criticism you open yourself up to.

Unfortunately, it’s hard for me to determine what “industrial strength”, “enterprise grade” and other such buzzwords are. But I’ll try to define high quality here, and try to show when a program is high quality.

## Organisation of this Document

I initially **wanted to give some examples** for open-source software that I considered to be of exceptional high-quality, but I decided against it.

That’s because their exceptional quality is a matter of taste, and it may provoke too much criticism against this essay as a whole. Thus, I’ll just give some examples, possibly accompanied by screenshots, of places where one program does something better than a different program. That does not imply that the latter is in fact of lower-quality than the former, in all possible respects.

This article is written from my point-of-view as a developer of “Free and Open Source Software” [[http://en.wikipedia.org/wiki/Free\\_and\\_open\\_source\\_software](http://en.wikipedia.org/wiki/Free_and_open_source_software)], and will focus on open-source and pseudo-open-source software. However, a lot of what I say is more universal.

## Parameters of Quality

This section will cover the parameters that make software high quality. However, it will not cover the means to make it so. These include such non-essential things as having good, modular code <sup>having\_good\_code</sup> good marketing, or having good automated tests. These things are definitely important, but will be covered only later, and a lot of popular, high-quality software lacks some of them, while its competition may do better in this respect.

One note that is in order is that these are **many parameters in a generic “weight function”**, and not a list of requirements which must all be satisfied. <sup>Licence</sup>

---

<sup>having\_good\_code</sup> What? Shouldn’t high quality software have a good codebase. Surprisingly no. What would you prefer: having a very modular codebase that does something pretty useless like outputting the string “Hello World” **or** having a large codebase of relatively low quality with a large amount of useful functionality and relatively few bugs?

I would certainly prefer the other alternative. That’s because I know I can always refactor that codebase, either by large scale refactoring [<http://www.joelonsoftware.com/articles/fog0000000348.html>] or by continuous refactoring or even “Just-in-Time” refactoring, only to add a new feature. <sup>Licence</sup> For example, I mention later that the more liberal the source licence of a program, the higher quality it is. Obviously, a lot of non- “Free and Open Source Software” [<http://en.wikipedia.org/wiki/FOSS>] or even binary-only applications are high-quality too. But, the use of such software is still more limited than an open-source one, so it would be of lesser quality than an identical software that is open-source.

## The Program is Available for Downloading or Buying

That may seem like a silly thing to say, but you'll be surprised how many times people get it wrong. How many times have you seen web-sites of software that claim that the new version of the program (or even the first) is currently under work, will change the world, but is not available yet? How many times have you heard of web-sites that are not live yet, and refuse to tell people exactly what they are about?

Alternatively, in the case of the "Stanford checker" [<http://en.wikipedia.org/wiki/Coverity>], which is a sophisticated tool for static code analysis, it is not available for download, but instead is a service provided by its parent company.

A program should be available in the wild somehow (for downloading or at least buying) so people can use it, play with it, become impressed or unimpressed, report bugs, and ask for new features. Otherwise it's just in-house software or at most a service, that is not adequate for most needs.

In the "Cathedral and the Bazaar", Eric Raymond recommends to "release early and release often" [<http://www.catb.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/ar01s04.html>]. Make frequent, incremental releases, so your project won't stagnate. If you take your project and work on it yourself for too long, people will give up.

If you have a new idea for a program, make sure you implement some basic but adequate functionality, and then release it so people can play with it, and learn about it. Most successful open-source projects, that have been open-source since their inception, have started this way: the Linux kernel, gcc, vim, perl, CPython. If you look at the earliest versions of all of them, you'll find that they were very limited if not downright hideous, but now they are often among the best of breed.

## The Version Number is Clearly Indicated

Which version of your software are you using? How can you tell? It's not unusual to come to a page where the link to the archive does not contain a version number, nor is it clearly indicated anywhere. What happens if this number was bumped to indicate a bug fix. How do you indicate it then?

A good software always indicates its version number in the archive file name, the opening directory file name, has a `--version` command-line flag, and the version mentioned in the about dialogue if there is one.

## The Source is Available, Preferably under a Usable Licence

Public availability of the source is a great advantage for a program to have. Many people will usually not take a look at it otherwise, for many reasons, some ideological but some also practical.<sup>Ideology</sup> Without the source, and without it being under a proper licence, the software will not become part of most distributions of Linux, the BSD operating systems, etc.

If you just say that your software "is copyright by John Doe. All Rights Reserved", then it may arguably induce an inability to study its internals (including to fix bugs or add features) without being restricted by a non-compete clause, or even that its use or re-distribution is restricted. Some software ship with extremely

---

Similarly, libraries or programs that are distributed under the relatively restrictive GNU General Public Licence (GPL) [[http://en.wikipedia.org/wiki/GNU\\_General\\_Public\\_License](http://en.wikipedia.org/wiki/GNU_General_Public_License)] (and which are considered open-source and usable by most Linux distributions) cannot be used in many common situations, and so would be of lesser quality than programs under a more permissive licence.

<sup>Ideology</sup> Assuming there was ever a true ideology that was not also practical.

The way I see it, an ideology (and ethics in general) are a strategy that aims to make a person lead a better, happier life. If it isn't, then it's just a destructive dogma, or just plain stubbornness.

long, complicated (and often not entirely enforceable) End-User-Licence-Agreements (EULAs) that no-one reads or cares to understand.

As a result, many people will find a program with a licence that is not 100% Free and Open Source Software [<http://www.gnu.org/philosophy/free-sw.html>] - unacceptable. To be truly useful the application also needs be GPL compatible [<http://www.dwheeler.com/essays/gpl-compatible.html>], and naturally usable public-domain licences such as the modified BSD licence, the MIT X11 licence, or even pure Public Domain source code<sup>public-domain</sup>, are even better for their ability to be sub-licensed and re-used. This is while licenses that allow incorporation but not sub-licensing, like the Lesser General Public License (LGPL) are somewhere in between.

While some programs on Linux have become popular despite being under non-optimal licences, many Linux distributions pride themselves that the core system consists of “100% free software”. Most software that became non-open-source, was eventually forked or got into disuse, or else suffered from a lot of bad publicity.

As a result, a high-quality software has a licence that is the most usable in the context of its common use cases. These licences are doubly-important for freely-distributed UNIX software.

## It “Just Works”

This is a meta-parameter for quality. When people say that something “just works”, they mean that you don’t have to be concerned about getting it up and running, not spend too much time to learn it, not worry about it destroying data, or have to wonder how to troubleshoot problems with it.

A program that just works is the holy grail of high-quality software. In practice this means several things:

1. A “just works” software also doesn’t have **any show-stopping bugs**. While it may still have some bugs, it should mostly function correctly.
2. It has most of the **features people want** and does not lack essential ones. For example, GNU Arch, an old and now mostly unused version control program, did not work on Windows 32-bit, while Subversion, a different and popular alternative, has a native port. Moreover, Mercurial, a different alternative, cannot keep empty directories (or trees of directories not containing files) in the repository. This may make both Mercurial and GNU Arch a no-starter for many uses.

Tendra [<http://www.tendra.org/>] is the most prominent alternative C and C++ compiler to GCC, but it’s hardly as advanced as GCC is, does not have all of GCC’s features and extensions, and is not usable as a replacement for GCC for most needs. As such, it is hardly ever used.

3. A “just works” software also has **good usability**. What it means is that it behaves like people expect it to. The Emacs-based editors, which are an alternative to Vim, do not invoke the menus upon pressing “Alt+F”, “Alt+E”, etc. which is the Windows convention to them.

Furthermore, when putting a single-line prompt, the prompt cannot be dismissed with either Ctrl +C or ESC, while in Vim, both keys dismiss the prompt. The key combination to dismiss it is not written anywhere on the screen and I won’t tell you what it is. According to User Interface Design for Programmers [<http://www.joelonsoftware.com/uibook/chapters/fog0000000057.html>], “A user-interface is well-designed when the program behaves exactly how the user thought it would”.

---

<sup>public-domain</sup> Using a pure-public-domain licensing terms for your software is problematic because not all countries have a concept of “public-domain”, similar to that of the United States, because many people misinterpret it, and because it is not clear whether software can be licensed under the public-domain to begin with. (And other such issues).

While quite a lot of important programs has been released under the public domain [<http://freshmeat.net/browse/197/>], and they are doing quite fine, they may have some problematic legal implications.

For these reasons, I now prefer the MIT X11 Licence for software that I originated instead of the “public domain”.

While some people may be led to believe this is not applicable to terminal applications, TTY applications, command line applications, or even Application Programming Interfaces (APIs) - it still holds there. One thing that made me like `gvim` (the Graphical front-end to `vim`) was that it could be configured to behave much like a Windows editor. I gradually learnt more and more `vim` paradigms, but found the intuitive usability a great advantage. But I could never quite get used to Emacs.

## The Program has a Homepage

Most software applications and libraries of high-quality have a homepage which introduces them, has download information, gives links, and provides a starting point to receive more information. And no - a `/project/myprogram/` page on Source Forge [<http://sourceforge.net/>] or a different software hub - is much more sub-optimal than that, and leaves a bad impression.

## The Program is Easy to Compile, Deploy and Install

A high-quality program is easy to compile, deploy and install. It builds out of the box with minimal hassles. There are several common standard building procedures for such software:

1. The common standard building procedure using the GNU Autotools [<http://sourceware.org/autobook/>] is `./configure --prefix=$PREFIX ; make ; make install`.
2. There are now some more modern alternatives to the GNU Autotools [<http://www.shlomifish.org/open-source/resources/software-tools/>], which may also prove useful.
3. CPAN [<http://www.cpan.org/>] Perl distributions have a similar `perl Makefile.PL` procedure or more recently also one using `perl Build.PL` which tends to be less quirky (see `Module-Build` [<http://cpan.uwinnipeg.ca/dist/Module-Build>]).

Generally, one usually installs them using the `CPAN.pm` or `CPANPLUS.pm` interfaces to CPAN, or preferably using a wrapper that converts every CPAN distribution to a native (or otherwise easy to remove) native system package.

4. Python [<http://www.python.org/>] packages have the standard `setup.py` procedure which can also generate Linux RPMs and other native packages.
5. There are similar building procedures for most other technologies out there.

However, it's not uncommon to find a program that fails to build even on GNU/Linux on an x86 computer, which is the most common platform for development. Or the case of the `qmail` email server [<http://www.shlomifish.org/open-source/anti/qmail/>], which has a long and quirky build process. It reportedly fails to compile on modern Linuxes, and someone I know who tried to build it said that it did not work after following all the steps.

One thing that detracts from a piece of software being high-quality is a large amount of dependencies.

If we take `Plagger` [<http://plagger.org/>], a web-feed mix-and-match framework in Perl (not unlike `Yahoo Pipes` [<http://pipes.yahoo.com/>], but predates it), then its `Plagger` distribution on CPAN [<http://search.cpan.org/dist/Plagger/>] contains all of its plug-ins inside, and as a result requires "half of CPAN" including such obscure modules, as those for handling Chinese and Japanese dates.

Popular programs like `GCC`, `perl 5`, `Vim`, `Subversion` and `Emacs` have very few dependencies and they are normally included in the package, if necessary to build the system. They are all written in very portable C and POSIX and have been successfully deployed on all modern UNIX-flavours, on Microsoft Windows and on many other more obscure systems.

While reducing the number of dependencies often means re-inventing wheels, it still increases the quality of your software. I'm not saying a program cannot be high-quality if it has a large amount of dependencies, but it's still a good idea to keep it to a minimum.

## The Program Has Packages for Most Common Distributions

A good program has packages for most common distributions, or such packages can be easily prepared.

Lack of such packages will require installing it from source, using generic binary packages, or other workarounds that are harder than a simple command to install the package from the package manager, and may prevent it from being maintained into the future.

A good example for how this can become wrong is the qmail SMTP server [<http://www.shlomifish.org/open-source/anti/qmail/>], before it became public-domain. The qmail copyright terms prevented distributing modified sources, or binary packages. As a result, the distributions that supported it packaged it as a source package, with an irregular build-process. Since the qmail package had its own unconventional idea of directory structure, some of the distributions had to extensively patch it. This in turn prevented more mainstream patches from being applied correctly to correct the many limitations that qmail had, or accumulated over the years due to its lack of maintenance.

## The Program Has Good, Extensive and Usable Documentation

If your GUI program is simple and well-designed, then you normally don't need good documentation. However, a command line program, a library, etc. does need one, or else the user won't know what to do.

There are many types of documentation: the `--help` flag [<http://freshmeat.net/articles/view/519/>], the man page, the README/USAGE/INSTALL files, full-fledged in-depth guides, documents about the philosophy, wikis, etc. If the program is well-designed, then the user should be able to get up and running quickly. An exception to this are various highly specialised programs, such as 3-D graphics programs, or CAD programs, that require some extensive learning.

If we take Subversion as an example, then it has a full Book online [<http://svnbook.red-bean.com/>], several tutorials, an `svn help` command which provides help to all the other commands, and a lot of help can be found using a Google search. GNU Arch, on the other hand, only had one wordy tutorial, that I didn't want to read. Most of the other tutorials people wrote, became misleading or non-functional as the program broke backwards compatibility.

Vim has an excellent internal documentation system. It's the first thing you are directed see when invoking it. It has a comprehensive tutorial, a full manual, and the ability to search for many keywords, with a lot of redundancy. As a result, one can easily become better and better with vim or gvim, albeit many people can happily use it with only the bare essentials.

Emacs' help on the other hand is confusing, dis-organised, lacking in explanation and idiosyncratic. It doesn't get invoked when pressing "F1", is not directed to when the program starts, and most people cannot make heads nor tails of it. There is a short Emacs tutorial, but it isn't as extensive as Vim's. Nor does it explain how to configure Emacs to behave in a better way than its default, in which it behaves completely differently to what people who are used to Windows-like conventions or vim-like conventions expect.

## Portability

It is tempting to believe that by writing a program for one platform, you can gain most of your market-share. However, people are using many platforms on many different CPU architectures: Windows 32-bit/64-

bit on Intel machines, Itanium, or x86-64; Linux on a multitude of platforms; BSD systems (NetBSD, FreeBSD, OpenBSD and others) on many architectures as well; Mac OS X; Sun Solaris (now also OpenSolaris), and more obscure (but still popular) Unix-clones like AIX, HP-UX, IRIX, SCO UNIX, Tru64 (formerly Digital Unix), etc. And to say nothing of more exotic, non-UNIX, non-Microsoft operating systems like Novell Netware, Digital Corp.'s VMS or OpenVMS, IBM's VM/CMS or OS/390 (MVS), BeOS, AmigaOS, Mac OS 9 or earlier, PalmOS, VxWorks, etc. etc.

As a general rule, the only thing that runs on top of all of these systems (in the modern “All the world is a VAX.” world) is a C-based program or something that is C-hosted.<sup>non-vax-like</sup> Most good programs are portable to at least Windows and most UNIXes and potentially portable to other platforms.

For example, Subversion has made it a high priority to work properly on Windows. On the other hand many of its early alternatives, especially GNU Arch, could not work there due to their architectures. As a result, many mixed shops, Windows-only shops, or companies where some developers wanted to use Windows as their desktop OS, could not use Arch. So Arch has seen a very small penetration.

The bootstrapping C compiler of gcc for example, is written in very portable K&R C [[http://en.wikipedia.org/wiki/C\\_\(programming\\_language\)#K.26R\\_C](http://en.wikipedia.org/wiki/C_(programming_language)#K.26R_C)], so it can be compiled by any C compiler. Later on, this compiler can be used to compile most of the rest of the GCC compilers.

Compare that to many compilers for other languages that are written in the same language. For example, GHC - The Glasgow Haskell Compiler [<http://www.haskell.org/ghc/>] is written in itself, and requires a relatively recent version of itself to compile itself. So you need to bootstrap several intermediate compilers to build it.

## Security

A high-quality program is secure. It has a relatively small number of security issues, and bugs are fixed there as soon as possible.

Some people believe that security is the most important aspect of software, but it's only one factor that affects its quality. For example, once I was talking with a certain UNIX expert, and he argued that the Win32 CreateProcess() system call [<http://msdn2.microsoft.com/en-us/library/ms682425.aspx>] was superior to the UNIX combination of Fork() and Exec() [<http://en.wikipedia.org/wiki/Fork-exec>], just because it made some bugs harder to code. However, some multitasking paradigms are not possible, without the fork() system call, which is not present in the Win32 API at all, and needs to be emulated (at a high run-time cost) or replaced with thread-based multitasking, which is not identical. Finally, it is still possible to get fork()+exec() right, and there's a spawn() abstraction on many modern UNIXes.

While I don't mean you shouldn't pay attention to security, or keep good security practices in mind when coding, I'm saying that it shouldn't slow down the process by much, or prevent too many exciting features from being added, or cause the development to stagnate.

## Backwards Compatibility

A high-quality program maintains as much backward compatibility [[http://en.wikipedia.org/wiki/Backward\\_compatibility](http://en.wikipedia.org/wiki/Backward_compatibility)] with its older versions as possible. Some backward compatibility, like relying

---

<sup>non-vax-like</sup> I'm fully aware that before C-based UNIX and UNIX-like systems became dominant there were some more exotic architectures that could not run C comfortably. Prime examples for them are the PDP-10 [<http://en.wikipedia.org/wiki/PDP-10>] and the Lisp machines [[http://en.wikipedia.org/wiki/Lisp\\_machine](http://en.wikipedia.org/wiki/Lisp_machine)].

However, such more “unconventional” architectures are now dead, and no CPU architecture developer in their right mind would want to create a CPU that won't be able to run C and C-based UNIX-based or UNIX-like operating systems such as Linux. (Unless it's probably a relatively niche micro-processor for embedded systems).

Lisp, and similar higher-level languages, run on modern UNIX-based OSes very well, so there's not a big problem there.

on bugs or other misbehaviours (“bug-to-bug compatibility”), is probably too extreme to consider. But users would like to upgrade the software and expect all of their programs to just continue to work.

A bad example for software that does not maintain backwards compatibility is PHP [<http://www.php.net/>], where every primary digit breaks the compatibility with the older one: PHP 4 was not compatible with PHP 3 and PHP 5 was not compatible with PHP 4. Furthermore, sometimes existing user-land code was broken in minor-digit releases. As such, maintaining PHP code into the future is a very costly process, especially if you want it to work with a range of versions.

On the other hand, the perl5 developers have been maintaining backwards compatibility between 5.000, 5.001, 5.002 up to 5.6.x, 5.8.x and now 5.10.x. Therefore, one can normally expect older scripts to just work. perl5 can also run a lot of Perl 4 code and below, and Perl 4 code can be ported to modern versions of perl5 with relative ease. While sometimes scripts, programs or modules were broken (due to lack of “bug-to-bug compatibility”), or became slower, upgrading to a new version of Perl is normally straightforward.

## Good Ways to Provide Support

A piece of high-quality software has good ways for its users to receive support. Some examples for ways to do that are:

1. A Mailing List.
2. IRC (Internet Relay Chats) [[http://en.wikipedia.org/wiki/Internet\\_Relay\\_Chat](http://en.wikipedia.org/wiki/Internet_Relay_Chat)] Channels.
3. An email address for questions.
4. Web Forums.
5. Wikis.

Without good ways to receive support, users will be unnecessarily frustrated when they encounter a problem, which cannot be answered by the documentation. Refer to Joel Spolsky’s “Seven Steps to Remarkable Customer Service” [<http://www.joelonsoftware.com/articles/customerservice.html>] for more information on how to give good support.

## Speed and High Performance

The reason I mentioned this quality parameter so late is because it was what Mr. Campbell stressed in his argument about “Industrial Strength” Freecell solvers. So I wanted to show that there are other important parameters beside it. However, raw performance is important, too.

If a program is too slow, or generates sub-optimal results, most people will be reluctant to use it and find using it daunting. They will either give up waiting for it to finish, or get distracted. If the output results of the program are too sub-optimal (assuming there’s a scale to their optimality), then they will probably look for different alternatives.

As a result, it is important that your software will run quickly, and will yield good results. There are many ways to make code write faster, and covering them here is out of the scope of this article.

A good example for how such optimisations can make such a huge difference are the memory optimisations done to Firefox between Firefox 2 and Firefox 3 [<http://blog.pavlov.net/2008/03/11/firefox-3-memory-usage/>], which greatly improved its performance, memory consumption, and reduced the number of memory leaks. It should be noted that often, reducing memory consumption can yield better performance because of a smaller number of cache misses, process memory swapping, and other such factors.

There are several related aspects of performance, that also affect the general quality and usability of a program. One of them is responsiveness, which is often manifested when people complain that the program is “sluggish”. Java programs are especially notorious for being such, for some reason, while programs written in Perl [<http://www.perl.org/>] and Python [<http://www.python.org/>] are more responsive and feel snappy, despite the fact that their backends are generally slower than the Java virtual machine.

A tangential aspect is that of startup time. Many programs require or have required a long time to start, which also makes using them frustrating, even if they are later responsive and quick.

## Aesthetics

A good program (or web site) or other resource is aesthetically pleasing. Aesthetics in this context, does not necessarily mean very “artsy” or having a breath-taking style. But we may have run into software (usually one for internal use or one of those very costly, bad-quality, niche, software) that seemed very ugly and badly designed, with a horrible user-interface, etc.

Different types of applications, and those running on different platforms, have different conventions for what is considered aesthetic. In *The Art of UNIX Programming* [<http://catb.org/~esr/writings/taoup/>], Eric Raymond makes the case for the “Silence is Golden” principle of designing UNIX command-line interfaces [<http://catb.org/~esr/writings/taoup/html/ch11s09.html>]. Basically, a command line program should output as little as possible. Now observe the behaviour of aptitude (a unified interface for package management) on Ubuntu Gutsy Gibbon, when trying to install a non-existing package name:

```
root@shlomif-desktop:/home/shlomif# aptitude install this-does-not-exist
Reading package lists... Done
Building dependency tree
Reading state information... Done
Reading extended state information
Initializing package states... Done
Building tag database... Done
Couldn't find any package whose name or description matched "this-does-not-exist"
The following packages have been kept back:
  firefox firefox-gnome-support
0 packages upgraded, 0 newly installed, 0 to remove and 2 not upgraded.
Need to get 0B of archives. After unpacking 0B will be used.
Reading package lists... Done
Building dependency tree
Reading state information... Done
Reading extended state information
Initializing package states... Done
Building tag database... Done
```

15 lines of output, and only one of them in the middle is the informative one. Why is all this information a concern of mine, especially given the fact that they are all given in the same monotonous default colour.

On the other hand, here's what urpmi (a similar package management interface for Mandriva) says on Mandriva Cooker:

```
[root@telaviv1 ~]# urpmi this-does-not-exist
```

```
No package named this-does-not-exist
[root@telaviv1 ~]#
```

Exactly one line and it's informative. While aptitude certainly has its merits, its verbosity still makes it much more painful to use than urpmi, when I have to work on Ubuntu.

Back to more visual aesthetics, one of the reasons that made me want to use Linux more than Windows 95' or 98' was the fact that its desktops were truly themable and could be made to look much better without effort. If I got tired of the same look, I could easily switch. While Windows XP shipped with a more attractive theme, and also had some proprietary and non-gratis theming software, Linux supplied all of that out-of-the-box and with a more attractive theme. The effects supplied by the Linux 3-D desktops [<http://youtube.com/watch?v=xC5uEe5OzNQ>], which have put the 9-billion Dollar effects of Vista to shame, have convinced some people to install Linux on their computer after seeing them.

## Conclusion

There are probably several parameters for software quality that I'm missing. However, the point is that one should evaluate the general quality of the software based on many parameters and not exclusively "security" or "speed" or whatever.

For example, many proponents of BSD operating systems claim that the various BSDs are superior to Linux because they are more secure, or because they are (supposedly) faster or are easier to manage, because their licence is less problematic than the GPL, etc. However, they forget that Linux has some advantages like being more popular (and so one can get support more easily), or like the fact that its kernel supports much more hardware, or that it has better vendor acceptance, and because more software is guaranteed to run with less problems on Linux than on the BSDs. <sup>linux-bsd-soft</sup>

I'm not saying the BSDs are completely inferior to Linux, just that Linux still has some cultural and technical advantages. Quality in software is not a linear metric, because it is affected by many parameters. If you're a software developer, you should aim to get as many of the parameters I mentioned right.

## Ways to Achieve Quality

Now that we've covered the elements that make programs high-quality here's a non-exhaustive list of ways to actually achieve it. None of them are absolutely necessary for achieving the quality, but they certainly help a lot.

Many people often confuse them with quality. "This software has God-damn awful code, so it's a pile-of-crap." Well, what do the users care how bad the code is, as long as it is functional, has all the necessary features, and is (mostly) bug-free? It's important not to confuse quality with the ways to achieve it.

The aim of this section is to briefly cover as many measures for achieving good quality as possible.

## Having Modular, Well-Written Code

The more modular a project's code is, the easier it is to change it, understand it and extend it, and the faster development will take. Refactoring [<http://www.refactoring.com/>] is the name given to the process used

---

<sup>linux-bsd-soft</sup> Naturally, this is a problem with the fact that most developers are developing on Linux (mostly x86), don't test it on other Unix flavours, and are too careless or unaware to write their programs portably enough.

However, it's still a quality parameter, because it still affects the way you're using the operating system.

to transform code from sub-optimal and “ugly”, but still mostly functional and bug-free code, to a code that is equally functional but more modular and clean. See the *Joel on Software* excellent article “Rub a dub dub” [<http://www.joelonsoftware.com/articles/fog0000000348.html>] for some of the motivation and practices of good refactoring instead of throwing away the code and restarting from scratch.

A reviewer of an early version of this article told me about an early and relatively large PHP code of her that was badly written, relatively buggy and yet proved to be popular among some of her clients, who have deployed it on many hosts and won’t effectively upgrade it. So she still has to maintain it, even though she’d rather not recommend it. She claimed that this is an indication that such low-quality in code, is a criterion of low-quality.

She has a point in the fact that usually badly-written or non-modular code results in more external low-quality factors such as bugs, security problems and lack of extensibility. However, even if the code was extremely well-written it is likely that it would need to be maintained, extended, and corrected. And if the clients in question don’t have or don’t want a good way to pull changes from a central place, or install updates properly, it’s a procedural and organisational problem.

Organisational quality deserves its own separate article (or arguably book, web-site, or even more than one book), but external software quality (much less internal one), is not a substitute for it. Please refer to a partial list I prepared on a different article [[http://www.shlomifish.org/philosophy/philosophy/advice-for-the-young/#recommended\\_writings](http://www.shlomifish.org/philosophy/philosophy/advice-for-the-young/#recommended_writings)], and my software “gurus” links [[http://www.shlomifish.org/links.html#software\\_gurus](http://www.shlomifish.org/links.html#software_gurus)] on my home-site’s links’ list.

## Automated Tests

Automated tests aim to test the behaviour of code automatically, instead of manual testing. The classical example for them is that if we wish to write a function called `add(i, j)` that aims to add two integers then we should check that `add(2, 3) == 5`, that `add(2, 0) == 2`, that `add(0, 2) == 2`, that `add(5, -2) == 3`, that `add(10, -24) == -14`, etc.

Then we can run all the tests and if any of them failed, we can fix them. Then after we write or modify the code, we can test using them again to see if there are any regressions.

Writing automated tests before we write the actual code, or before we fix a bug, and accumulating such tests (the so-called “Test-driven development” [[http://en.wikipedia.org/wiki/Test-driven\\_development](http://en.wikipedia.org/wiki/Test-driven_development)] paradigm), is a good practice which helps maintain high-quality code, and facilitates refactoring and makes it safer.

## Beta Testers

If you have beta testers for the code, or publish development versions frequently, you can get a lot of feedback for various different platforms and configurations your code is running on. These beta-testers can run the automated tests and also use the beta-code for their own testing or even production.

## Frequency of Releases

The more frequent your releases are, the more people can test your code, and the more they can upgrade to the latest version, and the quicker bugs that disturb your users are fixed, etc.

Naturally, there are advantages for slower release cycles, or for predictable release cycles like GNOME 2.x [<http://www.gnome.org/>] has. I won’t voice a definite opinion for which is the best methodology, but such a decision should be taken into consideration.

## Good Software Management

There are several sources, online and offline, explaining good software management for “shrinkwrap” software (open-source, commercial or other distributed) and for other types of software development (embedded, in-house, etc.), from which good advice can be taken for how to best run a software project. While they are sometimes contradictory, and often false, they still make a good read and are thought-provoking.

Here are some links:

1. Eric Raymond’s “The Cathedral and the Bazaar” series [<http://www.catb.org/~esr/writings/cathedral-bazaar/>]
2. The *Joel on Software* site [<http://www.joelonsoftware.com/>]
3. Paul Graham’s Essays [<http://www.paulgraham.com/>] (on various topics)
4. Extreme Programming [<http://www.extremeprogramming.org/>]

## Good Social Engineering Skills

It certainly helps for the project’s communities to have good social engineering skills. From tactfulness, to humour, to controlling one’s temper, to saying “Thanks” and congratulating people for their contribution, to timely applying patches and fixing bugs - all of these makes contributing to a project and using the program more fun and less frustrating.

Often, social engineering should be made part of the design of the software, or the web-sites dedicated to it. For example, it took me several iterations of having to fill the same project form in the GNU Savannah software-hub [<http://savannah.gnu.org/>], only for it to be rejected, and me having to follow the process again. Despite the fact the admins were polite, it still was annoying.

Eventually, they implemented a way to save previous project submissions and to re-send them, so future users won’t become as frustrated as I did.

Again, some projects have succeeded despite the fact they had, or even still have, bad social engineering. But adopting a good software engineering policy can certainly help a lot.

## No Bad Politics

Bad politics in a software project is a lot like subjectivity, it can never be fully eliminated, but we should still strive to reduce it to the minimum. [<http://shlomif.livejournal.com/52439.html>]. If bad political processes become common in a project, then important features are dropped, bugs are left unfixed, patches are stalled, external projects gets stalled or are killed, people become frustrated or angry and possibly leave the project - and the project may risk forking.

So it’s a good idea to keep bad politics at bay and to a minimum. How to do that is out of the scope of the document, but it’s usually up to the leaders to keep it so and maintain a good policy that will make as few people as possible frustrated and will not flabbergast external contributors. And naturally, for open-source and similar projects, a fork is often an option in this case, or in other similar cases.

## Good Communication Skills

A project leader and other participants should have good communication skills: very good English; pleasantness and tact; good phrasing, proper grammar, syntax, phrasing and capitalisation; clear writing;

patience and tolerance; etc. If they don't, then the project may encounter problems as people will find the project's developers hard to understand or tolerate, and, thus, hard to work with.

## Hype

Despite common belief I believe that the less hype and general noise there is regarding a software project, the better off it is. For example, as Paul Graham notes regarding Java [<http://www.paulgraham.com/javacover.html>]:

[Java] has been so energetically hyped. Real standards don't have to be promoted. No one had to promote C, or Unix, or HTML. A real standard tends to be already established by the time most people hear about it. On the hacker radar screen, Perl is as big as Java, or bigger, just on the strength of its own merits.

In fact, I can argue that if your project receives a lot of negative hype, then it is an indication that it is successful. Perl, for example, has received (and still receives) a lot of criticism [<http://perl-begin.org/learn/myth-dispelling/>], and Perl aficionados often get tired of constantly hearing or seeing the same repetitive and tired arguments by its opponents. However, the perl 5 interpreter is in good shape, it has many automated tests (much more than most other competing dynamic languages), an active community, many modules on CPAN ( the Comprehensive Perl Archive Network) with a lot of third-party , open-source functionality, and relatively few critical bugs. It is still heavily actively used and has many fans.

Similar criticism has been voiced against the Subversion version control system [<http://subversion.tigris.org/>], Linux, etc. One thing one can notice about such highly-criticised projects is that they tend not to be bothered by it too much. Rather, what they say is that "If you want to use a competing project, I won't stop you. It probably is good. It may be better in some respects. I like my own project and that's what I'm used to using and use."

This is by all means the right policy of "hyping" to adopt, if you want your project to be successful based on its own merits. Some projects compete for the same niche, without voicing too much hype against each other or for them, and this is a better indication that they are all healthy.

## A Good Name

Finally, a project should have a good name. One example for a project with an awful name is CVSNT [<http://www.march-hare.com/cvspro/>]. There are two problems with the name:

1. It is based on CVS, which most people have ran into its limitations, is considered passé and unloved, and people would rather avoid.
2. The "NT" part implies it only runs on Windows-NT, which is both misleading and undesirable.

On the other hand, the competing project "Subversion" has a much better name, since it has nothing to do with CVS, or Windows NT, and since it is an English word and sounds cool.

Some projects are successful despite being badly named, while some have a very cool name, but languish. Still, a good name helps a lot.

Also consider what Linus Torvalds said about Linux and 386BSD [<http://groups.google.com/group/comp.unix.pc-clone.32bit/msg/80bb74847934edc7>] (half jokingly):

No. That's it. The cool name, that is. We worked very hard on creating a name that would appeal to the majority of people, and it certainly paid off: thousands of people are using linux just to be able to say "OS/? Hah. I've got Linux. What a cool name".

386BSD made the mistake of putting a lot of numbers and weird abbreviations into the name, and is scaring away a lot of people just because it sounds too technical.

## Conclusion

Entire books (and web-sites) were filled with the various measures to achieve software quality, and what I described here was just a sample of it. The point is that these are not aspects of quality by themselves, but rather measures that help. None of them is a required or adequate condition for the success of a project, but the more are implemented the easier, faster, and more enjoying working on the project will be.

# Analysis of the Quality of the Various Freecell Solvers

## Introduction

As you recall, Gary Campbell's comment provided the motivation for my investigation into what makes a project high-quality and which quality-increasing-measures are not elements of quality by themselves. And since Freecell Solver [<http://fc-solve.shlomifish.org/>] has been my pet project, and as I'm still interested in techniques for solving Solitaire, I'd like to conclude this essay by analysing the various solvers according to the parameters I described. The natural caveat is that I may be somewhat biased due to the fact that I authored a solver of my own.

I was not so sure including this section is a good idea, but I feel the article is incomplete without it. Feel free to skip it, in case you are not that interested in it. <sup>FC-S-Comp</sup>

## The Analysis Itself

As is not uncommon in many software niches, some Freecell solvers are **not commonly available** for downloading or even for buying. For example, Bill Raymond's solver [<http://tech.groups.yahoo.com/group/fc-solve-discuss/message/158>] (named "Cat in the Sack") which he has been talking about extensively, has not been released yet, under any licensing terms. Similarly, Danny A. Jones has kept his solver (mentioned in the introduction) for himself, and has not released it to the wild in either source or binary forms.

As such, the utility of such solvers is heavily reduced. For one, they cannot be used as integrated solvers of Solitaire apps, because no one is going to wait for an email to the author to be sent, and the author to reply with the solution. Also, many researchers won't trust solutions given by them without the ability to inspect the source code.

Most other solvers [[http://fc-solve.shlomifish.org/links.html#other\\_solvers](http://fc-solve.shlomifish.org/links.html#other_solvers)] are **available for free download**, and many are accompanied with the source, most probably under an open-source licence.

---

<sup>FC-S-Comp</sup> This section is not meant as a comprehensive comparison of the Freecell solvers. The latter has yet to be written and would be much more difficult than what I'm doing here. It would be complicated by factors such as:

1. Some solvers are binary-only and only run on DOS or Windows. Open-source command-line ones may probably perform better on Linux or other Unix systems.
2. Some of them are not publicly available in any form.
3. Some of them can only solve Freecell, while others are more generic and can solve other Solitaire variants. (which in turn may hurt their performance.)
4. They differ in their solving algorithms. For example, some use atomic (= one-card) moves, some move entire sequences at a time, and some are based on meta-moves.

Gary Campbell's Solver [<http://www.numin8r.us/programs/>] is written in 8086 Assembly, with some 32-bit extensions, and only runs on DOS and DOS-compatible platforms. The assembly is compiled using a self-hosting macro assembler written by Campbell. The Assembly source code of the solver or the assembler are not available. As such modifying the code of the executable may prove to be problematic.

It is tempting to think that x86 and DOS are a very low common denominator, but that is not always the case. Imagine that you are writing a Freecell game for an embedded device running a non-x86 architecture such as ARM [[http://en.wikipedia.org/wiki/ARM\\_architecture](http://en.wikipedia.org/wiki/ARM_architecture)] or PowerPC [<http://en.wikipedia.org/wiki/PowerPC>]. In that case, in order to run Campbell's solver, you'll need to embed an x86-and-DOS emulator (such as DOSBox [<http://www.dosbox.com/>]) inside the device, which will complicate things, consume a lot of memory, and slow things down. On the other, a C-based solver, such as Freecell Solver, or Tom Holroyd's Patsolve, can be easily made to compile and run there with few if any modifications. So they would be preferable.

Similarly, when writing **portable programs** for the various Unix flavours [<http://en.wikipedia.org/wiki/Unix>] and other POSIX-capable or C-capable operating-systems, then writing code in x86 Assembly exclusively will make it a non-starter.<sup>Assembly</sup>

Some Freecell solvers also don't build out of the box. For example, typing make inside the patsolve distribution yields the following output:

```
shlomi:~/patsolve-3.0$ make
make clean
make[1]: Entering directory `/home/shlomi/patsolve-3.0'
rm -f patsolve *.o param.c param.h core win .depend
make[1]: Leaving directory `/home/shlomi/patsolve-3.0'
touch .depend
param.py param.dat
make: param.py: Command not found
make: *** [param.h] Error 127
```

This is relatively easy to fix, but still frustrating. Freecell Solver on the other hand, has been fully converted to the GNU Autotools, and can be built as a static and shared library (a.k.a DLL). It also has some built-in proof-of-concept, but still usable, command line utilities that link against it.

Freecell Solver has very good **usability**: it mostly respects the Unix conventions and best practices [<http://www.catb.org/~esr/writings/taoup/>], it can start solving from any arbitrary Solitaire board position given to it as input, and has kept command line interface usability in mind. Campbell's solver on the other hand, as of this writing, does not operate on a standard input/standard output manner, is poorly documented, and is counter-intuitive for someone who is used to Unix conventions (and most DOS conventions). I'm not even sure it can accept any arbitrary board as input, but I may be wrong.

Regarding **the homepage** of the solver: the Freecell Solver homepage [<http://fc-solve.shlomifish.org/>] has many pages, a navigation menu, a common look-and-feel and many links and information. On the other hand, Patsolve's homepage is nothing but a link in Tom Holroyd's software archive [[http://members.tripod.com/professor\\_tom/archives/](http://members.tripod.com/professor_tom/archives/)], which can easily be missed. Gary Campbell's homepage for his solver [<http://www.numin8r.us/programs/>] has a Baroque design and quite a lot of marketing-speak. And it's only one page with no anchors or a navigation menu, and very few links. Most other solvers don't fare better than that, and certainly worse than my own.

---

<sup>Assembly</sup> In some cases, the developers of portable software maintain versions of parts of the code in Assembly of selected architectures, while keeping more portable versions written in C or C++. This is done to optimise some parts for common CPU architectures.

**Licensing**, a necessary, but important, evil. The Freecell Solver's C source is distributed under the Public Domain licence, which allows virtually any use, including linking, modifying and sublicensing under any different licence by a third party.<sup>PD-LICENSE</sup> On the other hand, Patsolve is distributed under the GPL licence [[http://en.wikipedia.org/wiki/GNU\\_General\\_Public\\_License](http://en.wikipedia.org/wiki/GNU_General_Public_License)] which while usable is considerably more restrictive than Public Domain, or BSD-style licences. This is also the case for this Common Lisp solver [<http://kevin.atkinson.dhs.org/freecell/>] by Kevin Atkinson and Shari Holstege. Many other solvers, including Campbell's are binary-only, "All Rights Reserved", which makes them a complete no-starter for most open-source applications out there. Furthermore, many non-open-source applications will prefer to use a Public Domain solver, rather than paying royalties or risking "copyrights" and source availability problems.

Freecell Solver (FCS) is extensively documented, and even its online help is helpful and usable by itself [<http://freshmeat.net/articles/view/519/>]. The other solvers are much less documented, but arguably they also have much fewer options and features than Freecell Solver does.

Which brings us to the features - Freecell Solver has a long list of features [<http://fc-solve.shlomifish.org/features.html>], and I'm not aware of any other solver with so many. Especially of note is that it can solve many other Solitaire variants, including Simple Simon [[http://en.wikipedia.org/wiki/Simple\\_Simon\\_\(solitaire\)](http://en.wikipedia.org/wiki/Simple_Simon_(solitaire))], of which FCS is probably the only solver capable of solving.

Most of the solvers out there are fast, but some are more than others. Normally, for solving an individual game on the command-line, almost any solver will do, and the only cases where such speed matters more is for Solitaire research, and for analysing large sets of different games. (see the Freecell FAQ [<http://www.solitairelaboratory.com/fcfaq.html>] for more information).

In his original message, Mr. Campbell claimed that the Freecell Pro [<http://www.rhhistorical.com/rrdata/Fcpro65/>] solvers (Freecell Solver, Patsolve and the solver that originated from Don Woods) generated long and unusable solutions. That has not been my experience with Freecell Solver in some of its (practically infinite) configurations. Using the so-called "good-intentions" configuration I typically get solutions that are less than 200 moves. It is highly possible this is not the default in Freecell Pro. But like I said, there are many different parameters for quality than just speed and solution-length. Another fact worth of noting is that a beta-tester who tried out some of the Freecell Solver solutions in Freecell Pro, said that he found them to be very "creative" and interesting.

Freecell Solver has some other advantages: it is capable of being fully instantiated, as it stores everything in an "instance" C-struct, while using global variables only for constants (and not using static variables). This makes doing parallelised testing using mutli-threading much easier. It also has namespace-purity in the sense that all variables start with the `freecell_solver_` prefix. Moreover, FCS has a well-defined and stable API, which is not well-documented, but should be easy-to-use.

Part of the API is a parser for a list of command-line-like strings, which allow for configuring without making many standalone function calls.

## Conclusion

All things considered, I still feel that Freecell Solver is probably the best quality Solitaire solver out there. While it still has a lot of room for improvement [<http://fc-solve.shlomifish.org/to-do.html>], the rest of its competition have much bigger issues, or have not been made available (yet or ever).

---

<sup>PD-LICENSE</sup> As noted earlier, a Public Domain non-licence has some problems, which may make the software problematic for many corporations and in some jurisdictions. However, I don't have problem in exempting the licence of the Freecell Solver code (at least not the Public Domain code that I fully originated), from the Public Domain, including under the MIT X11 Licence, assuming this is necessary.

I'm also considering relicensing Freecell Solver and other older software of mine under the MIT X11 licence, or possibly having a dual-Public Domain and MIT X11 licensing terms (which seems somewhat silly, but may be a good idea.)

Several factors contributed for its success: the fact that I announced many releases on Freshmeat.net [<http://freshmeat.net/>], that I received a lot of input from my users and co-developers, that I was determined to constantly improve it and work on it, and that I worked on creating and maintaining a good web-site and documentation.

While most of the contributions of code I received were limited, the input from the users of the software proved to be crucial for its prosperity. Like I noted earlier, I lost interest in working on it (at least temporarily), but still maintain it, and feel that it is good enough as it is.

I hope this document, and similar resources it referenced will help you in working on your software, and improving its quality for the benefit of your users and you.

Happy Hacking!

## About This Document

### Copyrights

This work is licensed under the Creative Commons Attribution 2.5 License [<http://creativecommons.org/licenses/by/2.5/>] (or at your option a greater version of it). The CC-Attribution is almost Public Domain except for a requirement to make an attribution to the original author.

It was written by Shlomi Fish [<http://www.shlomifish.org/>] who also holds the copyrights.

### About the Author

Shlomi Fish was born in Israel in 1977, and has lived there most of his life. He is a user, developer, advocate and activist of Open Source Software. His single greatest contribution so far in this area has been Freecell Solver [<http://fc-solve.shlomifish.org/>], a Public Domain Library for solving games of Freecell and other types of Solitaire. However, he also initiated some other projects [<http://www.shlomifish.org/open-source/projects/>] and made important contributions to projects he did not initiate [<http://www.shlomifish.org/open-source/contributions/>] some of them very large scale. Recently, most of his contribution were perl distributions he uploaded to CPAN [<http://cpan.uwinnipeg.ca/~shlomif>] under his username.

Fish's work on Freecell Solver and other open source projects, and the fact he is interested in many software management techniques has led him to write a large number of articles [<http://www.shlomifish.org/philosophy/>], and weblog entries which describe his unique views as a developer of mostly portable software.

### Acknowledgements

Thanks to Mr. Gary Campbell, the author of "FCELL.COM" whose original message provided the inspiration for this article. Thanks to Omer Zak [<http://www.zak.co.il/>], who went over an early draft of this article and supplied many useful comments. Thanks to Jacinta Richardson of Perl Training Australia [<http://perltraining.com.au/>], who gave some very useful comments on the first and second revisions, and to "Limbic\_Region", for noting an omission from it. Thanks to many people on the IRC (= Internet Relay Chat) who have read the article and commented, and chose not to be explicitly credited.

Finally, thanks to all the past and present people who helped me with Freecell Solver, and with the rest of my Free and Open Source Software (FOSS) endeavours.