Photo: -JvL- (Flickr.com)

REVIEW

KUBUNTU 21.04

CREATE A CLONE
A HYBRID PORTABLE CLONE OF YOUR SYSTEM
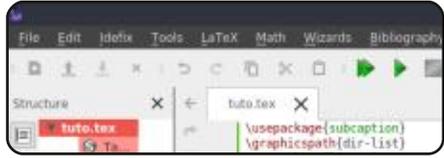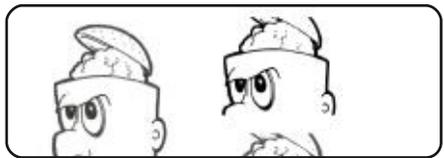
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

Command in the image:
```
#An alias to make the ls
command more detailed
alias ls = "ls -la --
color=always --classify"
```

EDITORIAL

WELCOME TO THE LATEST ISSUE OF FULL CIRCLE

Once again we have the usual suspects; Python, Latex, Inkscape, LMMS, MTMT (Micro This Micro That), and an article on creating a hybrid portable clone of your system. You never know when you'll need a copy of your system to boot from, or as a backup. Better safe than being in tears.

Richard returns to retrogaming in this month's Everyday Ubuntu, we continue our reviews of the 21.04 series with Kubuntu, an update to Ubports (Ubuntu) Touch in the form of OTA-18. Not to mention Q&A, a game review, and much more other goodness.

Don't forget: if you're looking for some help, advice, or just a chit chat: remember, we have a **Telegram** group. I mention this as I've been giving out the wrong link in a couple of emails. The (new and updated) link is: https://t.me/joinchat/24ec1oMFO1ZjZDc0. I hope to see you there. Come and say hello.

Anyway, stay safe, and all the best for 2021!
Ronnie
ronnie@fullcirclemagazine.org

FCM PATREON : https://www.patreon.com/fullcirclemagazine
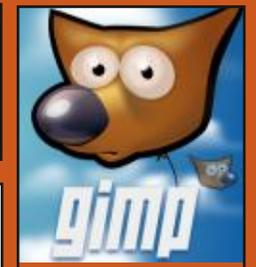
Find Full Circle on:

facebook.com/fullcirclemagazine

twitter.com/#!/fullcirclemag

http://issuu.com/fullcirclemagazine

https://mastodon.social/@fullcirclemagazine

Weekly News:

http://fullcirclemagazine.org/feed/podcast

http://www.stitcher.com/s?fid=85347&refid=stpr

http://tunein.com/radio/Full-Circle-Weekly-News-p855064/

LINUX 5.13 KERNEL RELEASE:

06/28/2021

After two months of development, Linus Torvalds has released the Linux 5.13 kernel. Among the most notable changes: initial support for Apple M1 chips, "misc" cgroup controller, discontinued support for /dev/kmem, support for new Intel and AMD GPUs, the ability to directly call kernel functions from BPF programs, randomization of the kernel stack for each system call, the ability to build in Clang with CFI (Control Flow Integrity) protection, LSM-module Landlock for additional limitation of processes, virtual sound device based on virtio, multi-shot mode in io_uring.

The new version received 17189 fixes from 2150 developers (the largest in history), the size of the patch is 60 MB (changes affected 12996 files, added 794705 lines of code, removed 399590 lines). About 47% of all changes introduced in 5.13 are related to device drivers, approximately 14% of changes are related to updating code specific to hardware architectures, 13% are related to the network stack, 5% are related to file systems and 4% are related to internal kernel subsystems.

https://lkml.org/lkml/2021/6/27/202

LTSM PROPOSED:

06/28/2021

The Linux Terminal Service Manager (LTSM) project has prepared a set of programs for organizing access to the desktop based on terminal sessions (so far using the VNC protocol). The developments of the project are distributed under the GPLv3 license.

Includes : LTSM_connector (VNC and RDP handler), LTSM_service (receives commands from LTSM_connector, starts login and Xvfb-based user sessions), LTSM_helper (graphical login interface), and LTSM_sessions (graphical interface for session management). The LTSM project differs from xrdp in other principles of work - the emphasis is on maximum use of built-in components, compact source code (less than a megabyte) and use of the DBus bus.

https://github.com/AndreyBarmaley/linux-terminal-service-manager

RELEASE OF MIXXX 2.3, THE FREE MUSIC MIXING APP:

06/29/2021

After two and a half years of development , the free Mixxx 2.3 package has been released , providing a complete set of tools for professional DJ work and creating music mixes. Downloads are available for Linux, Windows and macOS. The source code is distributed under the GPLv2 license.

http://mixxx.org/

UBUNTU IS MOVING AWAY FROM DARK HEADERS AND LIGHT BACKGROUNDS:

06/29/2021

In Ubuntu 21.10 approved the termination of deliveries of themes, combining dark headers, light background and light controls. Users will be presented with a completely light version of the Yaru theme by default , as well as the option to switch to a completely dark version (dark titles, dark backgrounds, and dark controls).

The decision is explained by the absence in GTK3 and GTK4 of the ability to define different background and text colors for the title and the main window, which does not guarantee the correct operation of all GTK applications when using combined themes (for example, in the gnome disk analyzer, a white input bar appears in the dark title). Another reason is too much labor costs to maintain non-standard themes. The problem

is that GNOME does not provide an official programming interface and set of guidelines for GTK themes, which leads to a break in compatibility with third-party themes in every new GNOME release.

Among other changes expected in Ubuntu 21.10, marked a departure from the use of aubergine color for the background of switches and widgets (color replacement has not yet been approved).

https://github.com/ubuntu/yaru/pull/2922

ULTIMAKER CURA 4.10 RELEASED:

06/29/2021

A new version of Ultimaker Cura is available, providing a graphical interface for preparing models for 3D printing (slicing). Based on the model, the program determines the scenario for the operation of a 3D printer with the sequential application of each layer. In the simplest case, it is enough to import the model in one of the supported formats (STL, OBJ, X3D, 3MF, BMP, GIF, JPG, PNG), select the speed, material and quality settings and send the print job. There are plugins for integration with SolidWorks, Siemens NX, Autodesk Inventor and other CAD systems. The CuraEngine engine is used to translate the 3D model into a set of 3D printer instructions . The project code is written in Python and is distributed under the LGPLv3 license. GUI built with frameworkUranium using Qt 5.

https://ultimaker.com/learn/an-improved-engineering-workflow-with-ultimaker-cura-4-10

POP!_OS 21.04 DISTRIBUTION OFFERS NEW COSMIC DESKTOP:

06/30/2021

System76 , a company specializing in the production of laptops, PCs and servers shipping with Linux, has released the Pop!_OS 21.04 . Pop!_OS is based on Ubuntu 21.04 and comes with its own COSMIC desktop environment. The project is distributed under the GPLv3 license. ISO images are generated for x86_64 architecture with NVIDIA (2.8 GB) and Intel / AMD (2.4 GB) graphics chips .

Prior to the release of Pop!_OS 21.04, the distribution came with a modified GNOME Shell, a set of original add-ons to GNOME Shell, its own theme, its own set of icons, other fonts (Fira and Roboto Slab), changed settings and an extended set of drivers. In the Pop!_OS 21.04 release, the modified GNOME desktop has been replaced by a new user environment, COSMIC (Computer Operating System Main Interface Components), which is developed under the GPLv3 license.

COSMIC continues to leverage GNOME technologies, but features conceptual changes and deeper desktop redesigns that go beyond additions to GNOME Shell. During the development of COSMIC, such tasks were set as the desire to make the desktop easier to use, expand functionality and increase work efficiency by customizing the environment to suit your preferences.

https://system76.com/pop

SEAMONKEY 2.53.8 INTEGRATED INTERNET APPLICATION SUITE RELEASED:

06/30/2021

SeaMonkey 2.53.8, which combines a web browser, an email client, a news feed aggregation system (RSS / Atom) and a Composer WYSIWYG editor for html pages within a single product. In the form of pre-installed add-ons, the Chatzilla IRC client, the DOM Inspector for web developers, and the Lightning calendar planner are offered. The new release brings fixes and changes from the current Firefox codebase (SeaMonkey 2.53 is based on the Firefox 60.8 browser engine with porting of security-related fixes and some improvements from the current Firefox branches).

https://www.seamonkey-project.org/news#2021-06-30

NEWS

SURICATA INTRUSION DETECTION SYSTEM UPDATE:

07/01/2021

The OISF (Open Information Security Foundation) has published corrective releases of Suricata Network Intrusion Detection and Prevention System 6.0.3 and 5.0.7 , which fix the critical vulnerability CVE-2021-35063. The problem makes it possible to bypass any Suricata analyzers and checks.

The vulnerability is caused by disabling flow parsing for packets with a nonzero ACK value but no ACK bit set, which allowed starting a TCP session with a SYN packet with a nonzero ACK to take the entire TCP connection out of scope in Suricata. Such packets in Suricata were recognized as erroneous and the handlers would return an error code without parsing the content.

https://suricata.io/2021/06/30/new-suricata-6-0-3-and-5-0-7-releases/

ALMALINUX INCLUDES SUPPORT FOR ARM64:

07/01/2021

The AlmaLinux 8.4, originally released for x86_64 systems, now supports the ARM / AArch64 architecture. There are three options of iso images available for download: bootable (650 MB), minimum (1.6 GB) and full (7 GB).

The distribution is fully binary compatible with Red Hat Enterprise Linux 8.4 and can be used as a transparent replacement for CentOS 8. Changes amount to rebranding, removal of RHEL-specific packages such as redhat-*, insights-client and subscription-manager-migration *, creation of a repository "devel" with additional packages and build dependencies. Includes new modules with Python 3.9, SWIG 4.0 Subversion 1.14, Redis 6, PostgreSQL 13 and MariaDB 10.5. Updated versions of GCC Toolset 10, LLVM Toolset 11.0.0, Rust Toolset 1.49.0 and Go Toolset 1.15.7.

https://wiki.almalinux.org/release-notes/8.4-arm.html

QUTEBROWSER 2.3 RELEASED:

07/01/2021

Qutebrowser 2.3, providing a minimal graphical interface that does not distract from viewing the content, and a navigation system in the style of Vim, built entirely on keyboard shortcuts, is out. The code is written in Python using PyQt5 and QtWebEngine. The source code is distributed under the GPLv3 license. The use of Python does not affect performance, since the rendering and parsing of content is carried out by the Blink engine and the Qt library.

https://lists.schokokeks.org/pipermail/qutebrowser-announce/2021-June/000104.html

TUX PAINT 0.9.26 IS RELEASED:

07/01/2021

The graphic editor for children - Tux Paint 0.9.26 has been published. The program is designed to teach drawing to children aged 3 to 12 years. Binaries are available for RHEL / Fedora, Android, Haiku, macOS and Windows. The latest release of the award-winning drawing program for children offers more fun for everyone with new features and expanded accessibility options.

http://www.tuxpaint.org/latest/tuxpaint-0.9.26-press-release.php

BACKDOOR DETECTED IN CLIENT SOFTWARE OF MONPASS CERTIFICATION AUTHORITY:

07/02/2021

Avast has published the results of a study of the compromise of the server of the Mongolian certification authority MonPass, which led to the substitution of a backdoor in the application offered for installation by customers. The analysis showed that the infrastructure was compromised by hacking one of MonPass's public web servers based on the Windows platform. Traces of eight different hacks were identified on the specified server, as a result of which eight webshells and backdoors for remote access were installed.

Including malicious changes were made to the official client software, which was shipped with the backdoor from February 8 to March 3. The story began when, in response to a client's complaint, Avast was convinced of the presence of malicious changes in the installer distributed through the official MonPass website. Upon notification of the issue, MonPass employees provided Avast with access to a copy of the compromised server's disk image to parse the incident.

https://decoded.avast.io/luigicamastra/backdoored-client-from-mongolian-ca-monpass/

JIM WHITEHURST, HEAD OF RED HAT, STEPS DOWN AS PRESIDENT OF IBM:

07/02/2021

Almost three years after the integration of Red Hat in IBM Vaytherst Jim has decided to withdraw from the post of president of IBM. At the same time, Jim expressed his readiness to continue to participate in the development of IBM's business, but already as an advisor to the IBM management. It is noteworthy that after the announcement of the departure of Jim Whitehurst, IBM shares fell in price by 4.6%.

From 2008 to 2019, Whitehurst served as chief executive officer (CEO) of Red Hat, and after completing the integration of Red Hat with IBM in 2019, he became vice president of IBM and head of the division into which Red Hat moved. In January 2020, Whitehurst was appointed by the board of directors to serve as president of IBM. Under Whitehurst's leadership, Red Hat's revenues have grown eightfold and market capitalization has increased tenfold.

https://www.cnbc.com/quotes/IBM

OPENZFS 2.1 RELEASE WITH DRAID SUPPORT

07/02/2021

The release of the OpenZFS 2.1 project has been published , which develops the ZFS file system implementation for Linux and FreeBSD. The project became known as "ZFS on Linux" and was previously limited to developing a module for the Linux kernel, but after the transfer of support for FreeBSD, it was recognized as the main implementation of OpenZFS and renamed.

The code is distributed under the free CDDL license, which is incompatible with GPLv2, which does not allow integrating OpenZFS into the upstream Linux kernel, since mixing code under the GPLv2 and CDDL licenses is not allowed. To work around this licensing incompatibility, it was decided to distribute the entire product under the CDDL as a separately downloadable module that is shipped separately from the kernel. The stability of the OpenZFS codebase is rated as comparable to other FS for Linux.

https://github.com/openzfs/zfs/releases/tag/zfs-2.1.0

NEOVIM 0.5, AVAILABLE:

07/03/2021

After nearly two years of development - Neovim 0.5 is out. This version focused on improving scalability and flexibility. The project has been overhauling the Vim code base for more than seven years, as a result - changes have been made to simplify the maintenance of the code, providing means of division of labor between several maintainers, separating the interface from the base part (the interface can be changed without touching the internals) and implementing a new extensible architecture for based on plugins. The original code of the project is distributed under the Apache 2.0 license, and the base part under the Vim license.

https://github.com/neovim/neovim/releases/tag/v0.5.0

AUDACITY'S NEW PRIVACY POLICY ALLOWS DATA COLLECTION FOR THE BENEFIT OF GOVERNMENT AUTHORITIES:

07/04/2021

Users of the Audacity sound editor have noticed a new privacy notice regulating issues related to sending telemetry and processing accumulated user information. In May, the audio

editor Audacity was sold to the Muse Group, which expressed its readiness to provide resources for modernizing the interface and implementing a non-destructive editing mode, keeping the product as a free project. Initially, Audacity was designed only to run on a local system, without accessing external services over the internet, but Muse Group plans to include tools in Audacity to integrate with cloud services, check for updates, send telemetry and reports with information about crashes and errors ... The Muse Group also tried to add code for accounting information about launching the application through Google and Yandex services, but after a wave of discontent, this change was canceled.

https://news.ycombinator.com/item?id=27724389

AᴮɪWᴏʀᴅ 3.0.5 ᴜᴘᴅᴀᴛᴇ:

07/04/2021

A year and a half after the last update, the release of the free multi-platform word processor AbiWord 3.0.5 was heralded. It supports the processing of documents in common office formats (ODF, OOXML, RTF, etc.) and provides features such as the organization of co-editing documents and multi-page mode. (This allows you to view and edit different pages of the document on one screen.) The project code is distributed under the GPLv2 license.

The new version fixes several bugs, including a crash when working with the clipboard. Fixed two vulnerabilities in the processor for MS Word formats, leading to a buffer overflow when processing specially designed footnotes and documents in the "doc" format.

http://www.abisource.com/release-notes/3.0.5.phtml

Tʜᴇ sᴇᴄᴏɴᴅ ᴇᴅɪᴛɪᴏɴ ᴏF ᴘᴀᴛᴄʜᴇs ғᴏʀ ᴛʜᴇ Lɪɴᴜх ᴋᴇʀɴᴇʟ ᴡɪᴛʜ sᴜᴘᴘᴏʀᴛ ғᴏʀ Rᴜsᴛ:

07/05/2021

Miguel Ojeda, author of the project Rust-for-Linux, proposed for consideration by the Linux kernel developers an updated version of the components for the development of device drivers on Rust language. Rust support is considered experimental, but has already been agreed for inclusion in the linux-next branch. The new version eliminates the comments made during the discussion of the first version of the patches. Linus Torvalds has already joined the discussion and suggested changing the logic for processing some bit operations.

https://lkml.org/lkml/2021/7/4/171

Rᴇʟᴇᴀsᴇ ᴏF Vɪʀᴛᴜᴏᴢᴢᴏ Lɪɴᴜх 8.4:

07/06/2021

Virtuozzo, who develops server virtualization software based on open source projects, published Virtuozzo Linux 8.4, constructed by recompiling the source code packages of Red Hat Enterprise the Linux 8.4. The distribution is fully binary compatible and identical in functionality to RHEL 8.4, and can be used to transparently replace solutions based on RHEL 8 and CentOS 8. Iso images of 1.6 GB and 4.2 GB are available for download .

https://www.virtuozzo.com/blog-review/details/blog/view/virtuozzo-vzlinux-84-now-available.html

OpenVMS operating system for x86-64 architecture:

07/06/2021

VMS Software, who bought the rights to continue development of the OpenVMS operating system from Hewlett-Packard five years ago, has released OpenVMS 9.1, notable for its support for the x86-64 architecture. The OpenVMS operating system has been evolving since 1977, is used in fault-tolerant systems that require increased reliability, and was previously only available for VAX, Alpha and Intel Itanium architectures. It is noteworthy that attempts to port OpenVMS to x86 systems in the late 1980s were unsuccessful due to the cessation of funding, after which the author of the port moved to Microsoft and created Windows NT.

https://vmssoftware.com/about/openvmsv9-1/

Nextcloud Hub 22 Collaboration Platform Available:

07/06/2021

Nextcloud Hub 22, which provides a self-contained solution for organizing collaboration between employees of enterprises and teams developing various projects, is out. Also, Nextcloud cloud platform 22, allowing the cloud to expand support synchronization and data exchange, providing the ability to view and modify data from any device at any point in the network (or using WebDAV web-interface). The Nextcloud server can be deployed on any hosting that supports PHP scripts and provides access to SQLite, MariaDB / MySQL, or PostgreSQL. The Nextcloud sources are distributed under the AGPL license.

https://nextcloud.com/blog/nextcloud-hub-22-introduces-approval-workflows-integrated-knowledge-management-and-decentralized-group-administration/

Tor Browser 10.5 released:

07.07.2021

After ten months of development, a major release of the specialized Tor Browser 10.5 is out, which continues to develop functionality based on the ESR branch of Firefox 78. The browser focuses on providing anonymity, security and privacy, all traffic is redirected only through the Tor network. It is not possible to contact the current system via the internet, which does not allow tracing the user's real IP. Tor Browser builds are prepared for Linux, Windows and macOS.

https://blog.torproject.org/new-release-tor-browser-105

Ubuntu 21.10 switches to using zstd algorithm for compressing deb packages:

07.07.2021

The Ubuntu developers began translating deb packages to use the zstd algorithm, which will almost double the speed of installing packages, at the cost of a small increase in their size (~ 6%). Notably, support for using zstd was added to apt and dpkg back in 2018 with the Ubuntu 18.04 release, but was not used to compress packages. On Debian, zstd support is already included in APT, debootstrap and reprepro and is being reviewed before being included in dpkg.

https://balintreczey.hu/blog/hello-zstd-compressed-debs-in-ubuntu/

Mozilla stops development of Firefox Lite browser:

07.07.2021

The Mozilla company decided to terminate the development of Firefox Lite, which is positioned as a lightweight version of Firefox the Focus, adapted to run on systems with limited resources and low-speed data links. The project was developed by the Mozilla development team from Taiwan and was aimed primarily at delivery in India, Indonesia, Thailand, the Philippines, China and developing countries.

Updates for Firefox Lite were discontinued on June 30th. Users are encouraged to switch to Firefox for Android instead of Firefox Lite . The reason for the termination of support for Firefox Lite is that in the current form Firefox for Android and Firefox Focus fully cover all the needs of mobile users, and the need to maintain another edition of Firefox has lost its meaning.

https://support.mozilla.org/en-US/kb/end-support-firefox-lite

NGINX 1.21.1 RELEASED:

07.07.2021

The release of the main branch of nginx 1.21.1 was announced,where development of new features continues (in the parallel maintained stable branch 1.20, only changes are made related to the elimination of serious bugs and vulnerabilities).

https://mailman.nginx.org/pipermail/nginx-announce/2021/000304.html

RELEASE OF PROXMOX VE 7.0:

07.07.2021

Proxmox Virtual Environment 7.0, a specialized Linux-distribution based on Debian, aimed at the deployment and maintenance of virtual servers using LXC and the KVM, and able to act as replacement products such as VMware vSphere, Microsoft Hyper-V and Citrix Hypervisor. The size of the installation iso image is 1 GB.

https://forum.proxmox.com/threads/proxmox-ve-7-0-released.92007/

SYSTEMD 249 SYSTEM MANAGER RELEASED:

07/08/2021

After three months of development, the system manager - systemd 249. The new release provides the ability to define users / groups in JSON format, stabilizes the Journal protocol, simplifies the organization of loading alternating disk partitions, adds the ability to bind BPF programs to services, implements mapping of user IDs in mounted partitions, offers a large portion of new network settings and the ability to run containers.

https://lists.freedesktop.org/archives/systemd-devel/2021-July/046672.html

RELEASE OF LINUX MINT 20.2:

07/08/2021

Linux Mint 20.2 distribution is out, based on the Ubuntu 20.04 LTS package base . The distribution is fully compatible with Ubuntu, but significantly differs in the approach to organizing the user interface and the selection of default applications. The Linux Mint developers provide a desktop environment that conforms to the classic canons of desktop organization, which is more familiar to users who do not accept the new methods of the GNOME 3 interface. DVD images: MATE 1.24 (2 GB) shells (2 GB), Cinnamon 5.0 (2 GB) and Xfce 4.16 (1.9 GB). Linux Mint 20 is classified as a Long Term Support (LTS) release, with updates to be rolled out until 2025.

http://blog.linuxmint.com/

STABLE RELEASE OF MARIADB 10.6 DBMS:

07/09/2021

After a year of development and three preliminary releases, the first stable release of the new MariaDB 10.6 DBMS branch is published. It is a fork from MySQL, while maintaining backward compatibility and featuring the integration of additional storage engines and advanced features. The new branch will be supported for 5 years, until July 2026.

MariaDB development is overseen by the independent MariaDB Foundation, following a completely open and transparent development process independent of individual vendors. MariaDB comes in place of MySQL on many Linux distributions (RHEL, SUSE, Fedora, openSUSE, Slackware, OpenMandriva, ROSA, Arch Linux, Debian) and has been deployed in major projects such as Wikipedia, Google Cloud SQL, and Nimbuzz .

https://mariadb.com/kb/en/mariadb-1063-release-notes/

SNOOP 1.3.0:

07/11/2021

The release of Project Snoop 1.3, which is an OSINT forensic tool that searches for user accounts in public data (open source intelligence), has been released. The program analyzes various sites, forums and social networks for the presence of the desired username, i.e. allows you to determine on which sites there is a user with the specified nickname. The project was developed on the basis of research work in the field of public data scraping. The downloads are prepared for Linux and Windows.

The code is written in Python and is licensed under a limited personal use license . At the same time, the project is a fork of the code base of the Sherlock project , supplied under the MIT license (the fork was created due to the inability to expand the base of sites).

https://github.com/snooppr/snoop/releases/tag/V1.3.0_10_July_2021

RELEASE OF EASYNAS 1.0 NETWORK STORAGE:

07/11/2021

EasyNAS 1.0, designed for the deployment of network attached storage (NAS, Network-Attached Storage) in small companies and home networks, is available. The project has been developing since 2013, built on the openSUSE package base and uses the Btrfs file system with the ability to expand the storage size without stopping work and creating snapshots. The size of the bootable iso image (x86_64) is 380MB. The 1.0 release marks a move to the openSUSE 15.3 package base.

https://easynas.org/2021/07/10/easynas-1-0/

GNU BINUTILS 2.37 RELEASED:

19/07/2021

The set of system utilities, GNU Binutils 2.37 is out, which includes programs such as GNU linker, GNU assembler, nm, objdump, strings, strip.

https://www.mail-archive.com/info-gnu@gnu.org/msg02923.html

HANDBRAKE 1.4.0 RELEASED:

19/07/2021

After almost two years if development, a release of a tool for multithreaded transcoding of video files from one format to another is presented - HandBrake 1.4.0. The program is available both in command line mode and a GUI interface. The project code is written in C (for Windows the GUI is implemented in .NET) and is distributed under the GPL license. Binaries for Linux (Flatpak), macOS and Windows are available for download.

Output, files can be generated in WebM, MP4 and MKV, AV1, H.265, H.264, MPEG-2, VP8, VP9 and Theora codecs. It can be used for video encoding, for sound - AAC, MP3 , AC-3, Flac, Vorbis and Opus. Additional functions include: bitrate calculator, preview during encoding, image resizing and scaling, subtitle integrator, wide range of conversion profiles for specified types of mobile devices.

https://handbrake.fr/news.php?article=47

LINUS TORVALDS COMMENTED ON THE SITUATION WITH THE NTFS DRIVER FROM PARAGON SOFTWARE:

19/07/2021

When discussing the issue of separation of powers when maintaining the code of file systems and VFS-related drivers, Linus Torvalds expressed his readiness to directly accept patches with the new implementation of the NTFS file system, if Paragon Software takes over the responsibilities of the NTFS file system maintainer in the Linux kernel and receives confirmation from others.

Linus noted that among the developers of the VFS kernel there are no people responsible for receiving pull requests with new filesystems, so such requests can be sent to him personally. In general, Linus hinted that he does not see any particular problems with the adoption of the new NTFS code into the main part of the kernel, since the deplorable state of the old NTFS driver does not

stand up to criticism, and no major claims were made to the new Paragon driver for a year.

The code for the new NTFS driver was discovered by Paragon Software in August last year and differs from the one already in the driver's kernel by the ability to work in write mode. The driver supports all the features of the current version of NTFS 3.1, including extended file attributes, data compression mode, efficient work with voids in files and replay of changes from the log to restore integrity after failures.

https://lore.kernel.org/lkml/CAHk-=wjW7Up3KD-2EqVg7+ca8Av0-rC5Kd7yK+=m6Dwk3D4Q+A@mail.gmail.com/

NVIDIA PUBLISHES DRIVER 470.57.02, OPENS RTXMU CODE AND ADDS LINUX SUPPORT TO RTX SDK:

19/07/2021

NVIDIA has published the first stable release of a new branch of the proprietary NVIDIA driver

470.57.02. The driver is available for Linux (ARM, x86_64), FreeBSD (x86_64) and Solaris (x86_64).

NVIDIA also released the code for the NVRHI (NVIDIA Rendering Hardware Interface) library and Donut framework under the MIT license. NVRHI is an abstract layer running on top of various graphics APIs (Direct3D 11, Direct3D 12, Vulkan 1.2) on Windows and Linux. Donut provides a set of pre-built components and rendering stages for prototyping real-time rendering systems.

In addition, NVIDIA has provided support for Linux and ARM architecture in the SDK: DLSS (Deep Learning Super Sampling, Realistic Image Scaling Using Machine Learning Methods), RTXDI (RTX Direct Illumination, Dynamic Lighting), RTXGI (RTX Global Illumination, Light Reflection), NRD (NVIDIA Optix AI-Acceleration Denoiser, Using Machine Learning to Accelerate Realistic Image Rendering).

https://forums.developer.nvidia.com/t/linux-solaris-and-freebsd-driver-470-57-02-production-branch-release/184049

FREETYPE 2.11 FONT ENGINE RELEASED:

19/07/2021

FreeType 2.11.0 , a modular font engine that provides a single API for unifying the processing and output of font data in various vector and raster formats.

https://www.mail-archive.com/freetype-announce@nongnu.org/msg00131.html

LINUX KERNEL BOOT IMPLEMENTED ON ESP32 BOARD:

20/07/2021

Enthusiasts were able to boot a Linux 5.0 environment on an ESP32 board with a Tensilica Xtensa dual-core processor (esp32 devkit v1 board, no full MMU) equipped with 2MB Flash and 8MB PSRAM connected via SPI. A ready-made Linux firmware image for ESP32 has been prepared for download.

The firmware is based on the

JuiceVm virtual machine and the Linux 5.0 kernel port. JuiceVm provides the smallest possible harness for RISC-V systems, capable of booting on chips with several hundred kilobytes of RAM. JuiceVm provides the launch of OpenSBI (RISC-V Supervisor Binary Interface), a glue interface for loading Linux kernel and minimal system environment from firmware specific to the ESP32 platform. Besides Linux, JuiceVm also supports FreeRTOS and RT-Thread booting .

https://www.reddit.com/r/esp32/comments/om106r/boot_linux_500_on_esp32/

Muse Group Seeks Closure of Musescore-Downloader Project GitHub Repository:

20/07/2021

The Muse Group, founded by the Ultimate Guitar project and owning the open source projects MusesCore and Audacity, has resumed attempts to close the musescore-downloader repository, that is developing an application for downloading music scores from musescore.com for free without having to log in to the site and without connecting to a paid Musescore subscription. (Pro.) The claim also concerns the musescore-dataset repository with a collection of sheet music copied from musescore.com. At the same time, Muse Group has nothing against the LibreScore project, in which the same author develops a free alternative to musescore.com, based on the code base of the MuseScore application, distributed under the GPL license. The position of the author of musescore-downloader boils down to the fact that he used a regular publicly documented API in his program, which was removed from musescore.com after the application was created. In addition, the author of musescore-downloader considers it incorrect that access to publications prepared by enthusiasts and initially placed in the open access free of charge was limited only to paid subscribers, while Muse Group does not own the rights to content prepared by users (users do not own the rights to the sheet music of other people's works, since the copyright holder is musicians and music publishers).

https://github.com/Xmader/musescore-downloader/issues/5#issuecomment-882450335

A GPIO driver written in Rust is proposed for the Linux kernel:

20/07/2021

In response to Linus Torvalds' comment that the driver example attached to the patchset for the Rust language support for the Linux kernel is useless and does not solve real-world problems, a version of the PL061 GPIO driver rewritten in Rust is proposed. A feature of the driver is that its implementation almost line by line repeats the existing GPIO driver in C. For developers who want to get acquainted with the creation of drivers in Rust, a line- by -line comparison has been prepared, which makes it possible to understand into which constructs in Rust the C code has been converted.

https://lore.kernel.org/ksummit/YPV7DTFBRN4UFMH1@google.com/

JuliaCon 2021 online conference will be held at the end of July:

20/07/2021

From July 28 to July 30, the annual JuliaCon 2021 conference will be held. Dedicated to the use of the Julia language and designed to perform high-performance scientific computing. This year the conference will be held online, registration is free.

From today until July 27, a series of thematic seminars will be held for the conference participants, where the solution of specific problems will be considered in detail. Seminars suggest a different level of familiarity with the language: from zero to advanced.

The upcoming conference will cover both programming issues and solutions to various applied problems from various fields, including those not related to scientific research. The conference is aimed both at participants who are just getting acquainted with the possibilities of the language and at

advanced users.

https://juliacon.org/2021/

STOCKFISH SUES CHESSBASE AND REVOKES GPL LICENSE:

20/07/2021

Project Stockfish, distributed under the GPLv3 license, filed a lawsuit on ChessBase. Stockfish is the strongest chess engine used by the lichess.org and chess.com chess service . The lawsuit is for the inclusion of Stockfish code in a proprietary product without opening the source of a derivative work.

The company ChessBase known for its chess program Fritz from the 1990s. In 2019, they released the Fat Fritz engine, based on the neural network of the open source Leela Chess Zero engine, at one time based on the AlphaZero project, opened by Google. This was not a violation of any law, although the Leela developers were unhappy that ChessBase positioned Fat Fritz as a standalone development, not recognizing the AlphaZero and LeelaZero teams.

As a reminder, the GPL license provides for the possibility of revoking the license from the violator and terminating all the rights of the licensee granted by this license. In accordance with the rules of license termination adopted in GPLv3, if violations were detected for the first time and eliminated within 30 days from the date of notification, the rights to the license are restored, and the license is not completely revoked (the contract remains intact). The rights are returned immediately also in case of elimination of violations, if the copyright holder has not notified about the violation within 60 days. If the deadlines have expired, then the violation of the license can be interpreted as a violation of the contract, for which financial penalties can be obtained from the court.

https://stockfishchess.org/blog/2021/our-lawsuit-against-chessbase/

VIRTUALBOX 6.1.24 RELEASED:

20/07/2021

Oracle has published a patch release for VirtualBox 6.1.24 , which includes 18 fixes .

https://www.mail-archive.com/vbox-announce@virtualbox.org/msg00212.html

LINUX KERNEL ROOT VULNERABILITY AND SYSTEMD DENIAL OF SERVICE:

21/07/2021

Security researchers at Qualys have disclosed details of two vulnerabilities affecting the Linux kernel and the systemd system manager. A kernel vulnerability (CVE-2021-33909) allows a local user to achieve root code execution by manipulating heavily nested directories.

The vulnerability is compounded by the fact that the researchers were able to prepare working exploits that work in Ubuntu 20.04 /

20.10 / 21.04, Debian 11 and Fedora 34 in the default configuration. It is noted that other distributions have not been tested, but theoretically they are also susceptible to the problem and can be attacked. The full code of the exploits is promised to be published after the widespread elimination of the problem, and so far only a prototype of limited functionality is available, causing the system to crash. The issue has been apparent since July 2014 and affects kernel releases since 3.16. The vulnerability patch was coordinated with the community and accepted into the kernel on July 19. Major distributions have already generated kernel package updates (Debian ,Ubuntu , Fedora , RHEL , SUSE , Arch).

https://blog.qualys.com/vulnerabilities-threat-research/2021/07/20/sequoia-a-local-privilege-escalation-vulnerability-in-linuxs-filesystem-layer-cve-2021-33909

KDE Plasma Mobile 21.07 Released:

22/07/2021

KDE Plasma Mobile 21.07 is published, based on the mobile edition of the Plasma 5 desktop, the KDE Frameworks 5 libraries, the Ofono phone stack and the Telepathy communication framework. Composite server kwin_wayland is used to display graphics. PulseAudio is used for sound processing.

The structure includes applications such as KDE Connect to pair your phone with a desktop, document viewer Okular, music player VVave , image viewer Koko and Pix, system reference notes Buho, calendar planner calindori, File Manager Index, the application manager Discover , program SMS sending Spacebar, plasma-phonebook, plasma-dialer , plasma-angelfish browser and Spectral messenger...

https://www.plasma-mobile.org/2021/07/20/plasma-mobile-gear-21-07/

SixtyFPS 0.1.0 GUI library available:

22/07/2021

The cross-platform library for creating graphic interfaces SixtyFPS 0.1.0 has been published, targeted for use on embedded devices and desktop applications on Linux, macOS and Windows platforms, as well as for use in web browsers (WebAssembly). The library code is written in Rust and is distributed under the GPLv3 or commercial license, which allows use in proprietary products without opening their code.

The authors of the library (Olivier Goffart and Simon Hausmann), former KDE developers who later moved to Trolltech to work on Qt, have now founded their own company developing SixtyFPS. One of the goals of the project is to provide the ability to work with minimal consumption of CPU and memory resources (several hundred kilobytes of RAM are required to work). There are two backends available for rendering - gl based on OpenGL ES 2.0 and qt using Qt QStyle.

https://github.com/sixtyfpsui/sixtyfps/releases/tag/v0.1.0

Firewalld 1.0 released:

23/07/2021

Introducing firewalld 1.0, implemented in the form of a binding over nftables and iptables packet filters. Firewalld runs as a background process that allows you to dynamically change the packet filter rules over D-Bus, without having to reload the packet filter rules and without breaking any established connections. The project is already used in many Linux distributions, including RHEL 7+, Fedora 18+, and SUSE / openSUSE 15+. The firewalld code is written in Python and is licensed under the GPLv2 license.

A major change to the version number is due to changes that break backward compatibility and change the behavior of working with zones. All filtering parameters defined in the zone are now applied only to traffic addressed to the host on which firewalld is running, and to filter transit traffic, you need to configure policies .

https://firewalld.org/2021/07/firewalld-1-0-0-release

Gentoo started generating additional assemblies based on Musl and systemd:

23/07/2021

The developers of the Gentoo distribution have announced that they have expanded the range of ready-made stagefiles available for download. The publication of stage-archives based on the Musl C library and assemblies for the ppc64 platform, optimized for POWER9 processors, has begun. They added stage archives with systemd manager for all supported platforms, in addition to previously available builds based on OpenRC. Through the standard download page for the amd64 platform, the delivery of Hardened stage files with SELinux support and the musl library began.

The changes were made possible by the introduction of new build hosts. Assemblies for amd64, x86, arm (via QEMU) and riscv (via QEMU) architectures are now

generated on a server with an 8-core AMD Ryzen 7 3700X CPU and 64 GB of RAM. Builds for ppc, ppc64, and ppc64le / power9le architectures are provided on a server with a 16-core POWER9 CPU and 32GB of RAM. For arm64 assemblies, a server with an 80-core Ampere Altra CPU and 256 GB of RAM is dedicated.

https://www.gentoo.org/news/2021/07/20/more-downloads.html

PeerTube 3.3 decentralized video broadcasting platform released:

25/07/2021

PeerTube offers a vendor-independent alternative to YouTube, Dailymotion and Vimeo, using a content distribution network based on P2P communications and linking visitors' browsers. The project ode distributed under the AGPLv3 license.

https://joinpeertube.org/news#release-3.3

Hardware Health Assessment Initiative in a Future Release of Debian 11:

23/07/2021

The community has launched an open beta test for the upcoming release of Debian 11, in which even the most inexperienced novice users can participate. Full automation was achieved after the hw-probe package was included in the new version of the distribution, which can independently determine the operability of individual devices from the logs.

A daily updated repository with a list and catalog of tested hardware configurations was organised. The repository will be updated until the expected release of a new version of Debian on August 14, after which it will become the place for collecting system usage statistics for the next two years.

https://github.com/linuxhw/TestCoverage/tree/main/Dist/Debian_11

The VirtualBox Networking Primer
Connecting and Configuring Virtual Machines

The VirtualBox Networking Primer is a no-nonsense guide for the VirtualBox user taking their next steps into virtual networks.

While Oracle VM VirtualBox is a great free tool, the real power of virtualisation comes when you start connecting virtual machines to each other and to the wider world. Software development, sales, education and training are just some of the areas in which network access to virtual machines offers endless opportunities. But the world of computer networks is filled with complex technical jargon.

Complete with principles, practice, examples and glossary, The VirtualBox Networking Primer takes the frustration and confusion out of connecting real-world projects.

Author: **Robin Catling**
Publisher: Proactivity Press
ISBN13 : 9781916119482

Amazon US link:
https://www.amazon.com/dp/1916119484?ref_=pe_3052080_397514860

Amazon UK link:
https://www.amazon.co.uk/VirtualBox-Networking-Primer-Connecting-Configuring/dp/1916119484/ref=sr_1_1?dchild=1&keywords=virtualbox+networking+primer&qid=1600253699&s=books&sr=1-1

Kobo:
https://www.kobo.com/us/en/ebook/the-virtualbox-networking-primer

I hope you've paid attention to previous issues, as this month we will be steaming ahead. If you did not, I strongly advise that you do. However, to avoid confusion, we are looking into ZynaddsubFX (Zyn from now on) that comes with LMMS. I want you to know that, as ugly as Zyn is (ugly as Zyn... OK, I'll see myself out), you can make almost any synth out there on the market with this. Not kidding!

Fire up LMMS and let's start!

First, Addsynth. Pull Zyn into your song editor. Click the show GUI button. Click advanced. Click edit instrument. Click the add button under addsynth. Click show voice parameters. Next to the sine wave, click change.

This may look like an equalizer, but it is not. It is a little counter-intuitive, the base function is to the left and the Oscillator is to the right, but you get used to it.

In the equalizer-type board that fills the bottom half of the screen, the even numbers usually mean an octave (just keep that in mind). The first one, top and bottom is a special one. Use this to make your instruments sound 'fuller'. You should experiment with even numbers as well as 2, 4, 8, 16, 32.

In the Addsynth voice parameters, one screen back, you will find your new voice / instrument in the picture window. Just below that in the bottom-left, you should see "Unison". It is off by default. You can add as many as you like, and by using the slider right next to it, you open the "Frequency Spread" that allows Zyn to move the tuning a bit – giving you a wider spectrum of sound, if you would.

Please, you have to play with this and hear the difference to understand; I cannot tell you; you need to learn this, so if you do not follow along and fiddle (don't worry, you won't break it) you will not understand what it sounds like. This will allow you to turn a simple chord into a trance piano sound, for instance, if you ever wondered where they get those sounds.

Then next to the "Frequency Spread" you should see a knob with the word "Stereo" above it; this spreads the sound in the left / right spectrum. The more you crank this knob, the 'wider' the sound will be. Try it! Spreading your instruments in stereo allows for better sounding tunes in the end.

If your goal is to create special effects for your games, the next two knobs are your friend. Vibrato makes those warbling sounds. Finer tuning for these can be done in the frequency LFO, but again, I cannot tell you, you need to actually play with it and find what you like or what works for you. If you have a large screen, move the piano out to the side, as there is no "play" option here. If you make a boo-boo, just click on clear to the right, to reset all the harmonics.

The nice thing about this is that each "instrument" you create has its own volume control, so you can add loud and soft versions of the same instrument to use in your song, without having to fiddle with volume automation. You can also
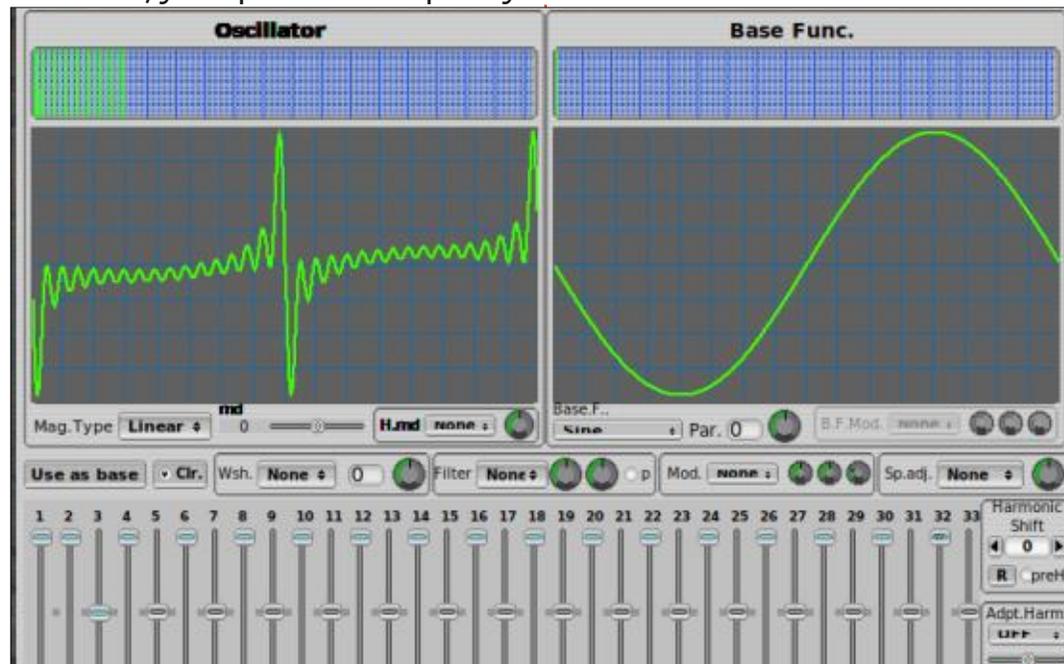
have multiple instances of Zyn in your song editor, so don't be shy.

Also remember that all these are dependent on your sound card too. My laptop's built-in sound card is probably made of cardboard, so it is quite horrible. When I connect my external Soundblaster card, the difference is chalk and cheese. Thus, if you are just mucking about like we are, don't try to make anything too serious, you need a decent sound card for that.

On to Padsynth, (I am skipping Subsynth as I don't see how it is useful.)

Things may look different and also similar, but it will sound different than Addsynth. If you look at the interface, you will see 'apply changes' is highlighted in a different color. (Supposedly red) as you need to press this every time to apply the changes you have made, before you will hear them. You would be surprised how often you forget.

To the right of the sine wave, we have our change button, which will allow us to choose a new wave type, the same as in the other two synths. The reason the Padsynth

sounds so much different to the Addsynth is – if you look above the wave, you will see two blue dotted lines intersecting the cyan harmonic. This is adjustable by the bandwidth knob in the bottom-left. It tells you above that that it is the frequency distribution. That is what is filling out the horrible PC speaker-esque sound to something that does not make your ears bleed.

In the image, top-left, there is a word that is marred by the outline, just below 'base type'. It says Gauss; changing that to square should get you into the synthwave-type sounds. I know some of you have

been on a synthwave trip (shout out to Peter), now here is your chance to create your own! The next one on that list, double exp, is basically a combination of the two above. You can also use the "AmpMultiplier" to multiply the amp, but what does that mean? It creates another hidden copy of the harmonic and allows you to shape the tune a bit more. Again, if you have tinny speakers, you may not hear the difference, so I urge you to get semi-decent speakers and a semi-decent sound card if you plan on playing with Zyn. Or… if you are tone deaf like me, you can just sputter ahead. (No, really, I listened to Rob Lande's video of violins that


cost between $70 - $10 million and they all sounded the same to me!).

As the name says, it is made for making pads, they sound really great when you use octaves? (bear with me on the music terminology, I don't know it, I know what I like and I know how to get it, and that is enough for me). Bear in mind that this is a base. Top-right it will even say "base type".

I have led you to the 'points of interest' – however, the onus is on YOU to fiddle and find sounds you like. This is a really great way to do remixes of songs with arb instruments.

Hope you guys had fun; we will continue this in the next issue.

As always, redirect your queries to misc@fullcirclemagazine.org

Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.
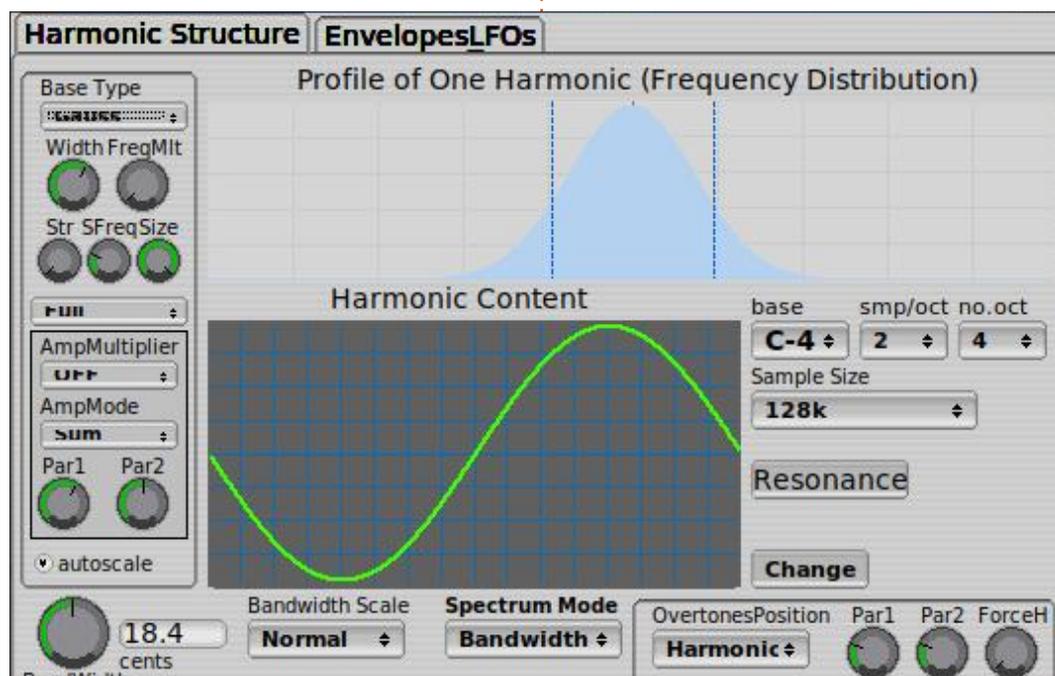
In the poem "to a mouse", Robert Burns wrote "*The best laid schemes o' Mice an' Men Gang aft agley*". Now, some of us might not get Scottish quite readily, so with a little bit of web searching, the general gist of the saying is (and this is the opinion of multiple website writers) "*the best laid plans of mice and men can still go wrong. No matter how well you plan, there is no guarantee of success.*" That seems to be speaking directly to my life.

The article you are reading was not my first choice of subjects for this month. In fact, it didn't make it into the top three subjects. But such is life. When life gives you lemons, what do you do? You make Lemon Drop Martinis. That's what.

PLOTEXT

I'm always looking for a way to produce plots that I can use easily, quickly and differently. I found what I think is a really good one. 'Plotext' plots directly on terminal; it has no dependencies; and the syntax is very similar to matplotlib.

It also provides a simple command-line tool.

I haven't had a tremendous amount of time to fully test and examine plotext, but by the time you get to the end of this article, you will see that it isn't very difficult to use.

You can see the project at https://github.com/piccolomo/plotext .

It's easy to install using pip (or pip3).

`pip install plotext`

Once I had the library loaded, I pulled open a ptpython REPL. I borrowed one of the many examples from the repository (shown right).

WORKING WITH REAL DATA

I wanted to see just how easy it is using some real data. I did a quick web search for the latest Covid-19 data for confirmed cases in Texas from March 3, 2020 until July 12,

```
import plotext as plt
y1 = plt.sin(1000, 3)
y2 = plt.sin(1000, 3, 1.5, phase = 1)
plt.plot(y1, label = "plot")
plt.scatter(y2, label = "scatter", marker = "small")
plt.plotsize(100, 30)
plt.title("Multiple Data Set")
plt.show()
```
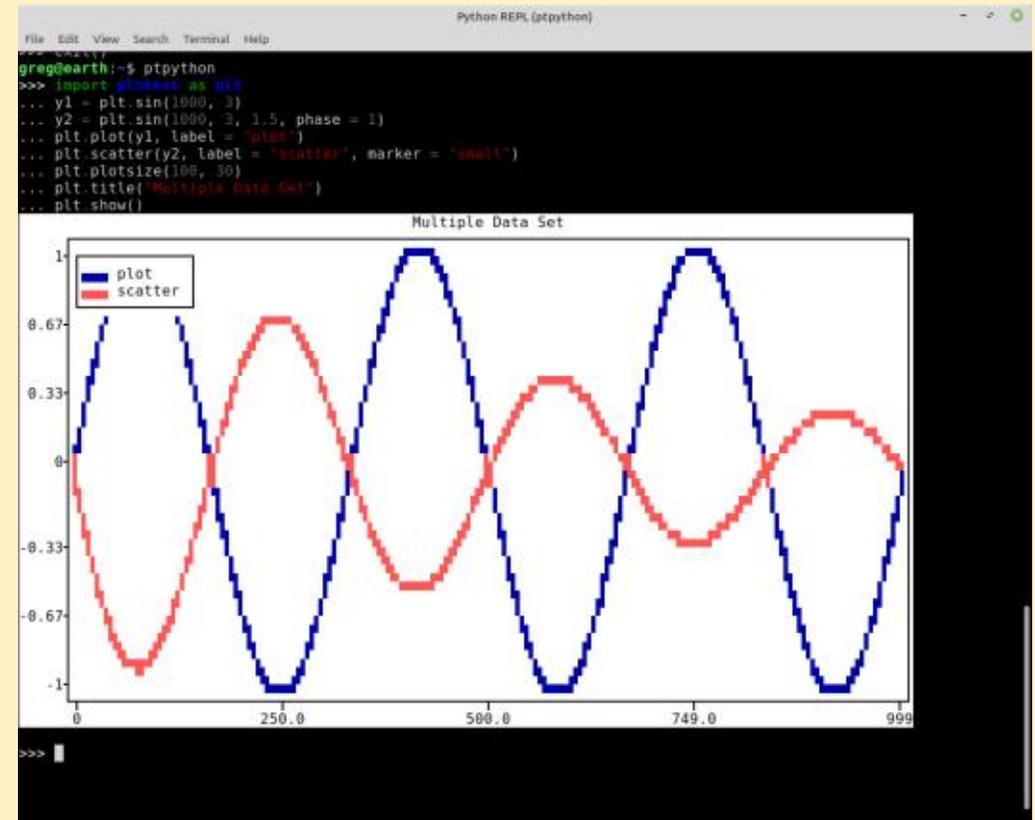

2021. It was a very large amount of data, so I simply pulled up the data in a spreadsheet and pulled the totals for all of Texas for the last 40 days. Then I threw together the following program (next page, top right).

You can see that the 'last 40 days' shows what the news is reporting. Despite the attempts to control Covid-19, the number of confirmed cases rose at a fairly steep rate.

Next, I tried to pull the data just for one county in Texas (bottom right). Harris county is where most of Houston is located. Houston has over 4 million residents and is the fourth largest city in the United States.

This time I pulled only 28 days worth of data. This plot shows that, in Harris county, things were starting to stabilize and actually going down until around the 28th of June, then started going up at a fairly quick rate.

Finally, I pulled the data for Hays County in Texas (next page, top right). Hays county is a fairly small county in terms of population, only around 157,000. Many of the residents of Hays county are college students.

Again, the number of confirmed cases went from 17150 to 19230 in a period of 103 days. While that's only 2088 more cases over the sample period, I find it interesting

```
import plotext as plt

data = [
    2520151, 2522016, 2523617, 2524212, 2524688, 2525991, 2527350, 2529068,
    2530176, 2530658, 2530910, 2531194, 2532112, 2533123, 2534149, 2534730,
    2536217, 2536646, 2536691, 2538120, 2538908, 2540318, 2541586, 2542877,
    2543158, 2543253, 2544966, 2546040, 2547411, 2548878, 2549911, 2550317,
    2550615, 2551284, 2552598, 2555504, 2557363, 2559192, 2559734, 2560137
]
print(data)
plt.plot(data)
plt.plotsize(100, 40)
plt.title('Covid-19 numbers last 40 days')
plt.show()
```


When I ran the program, this is the plot that I got:
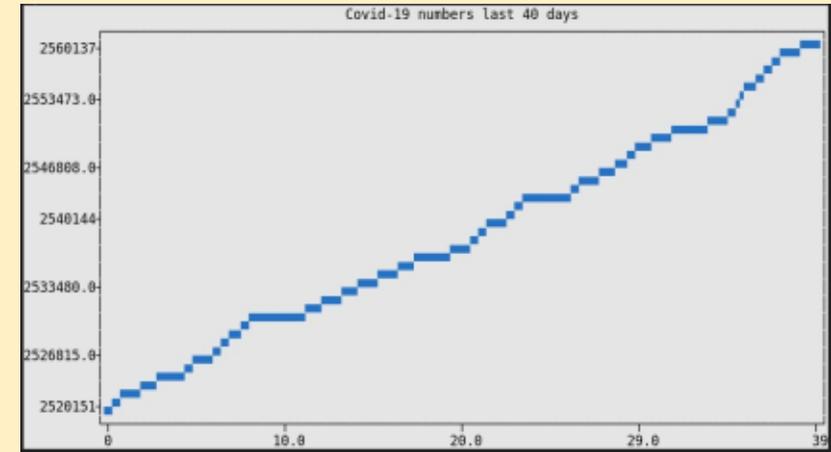
```
import plotext as plt
data1 =
[403401,403496,403151,402894,402916,403069,402988,403035,403187,403027,403059,403134,
403276,403149,403149,402927,403019,403205,403240,403512,403584,403645,403646,404402,4
04986,405327,405627,405757]
plt.plot(data1, label='Harris County')
plt.plotsize(100,30)
plt.title("Harris County")
plt.show()
```


When I ran the program, this is the plot that I got:
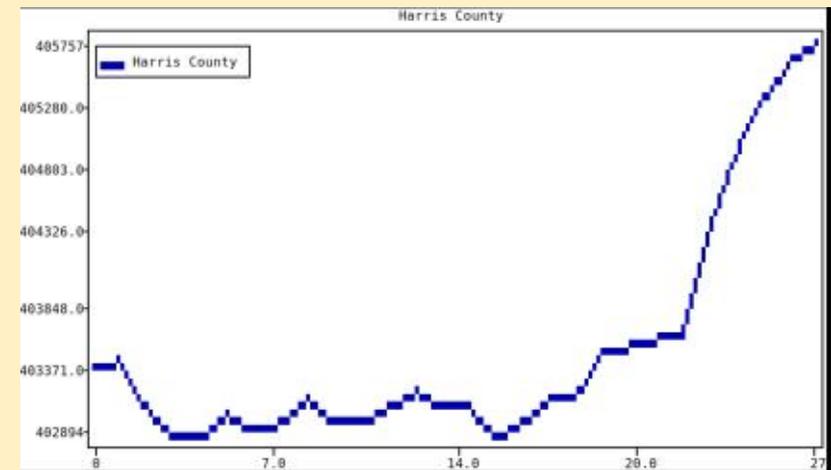
that the number of cases were rising at a fairly steady rate until about the time that the students started heading home for the summer break. Then the number of cases still went up but at a much slower rate.

As you can see, with a very little amount of effort, you can get reasonable plots. According to the repository site, you can do scatter plots, line plots, log plots, stem plots, multiple data sets, double Y axis plots, bar plots, histogram plots, set the data ticks to custom values, date/time plots, multiple subplots, plot streaming data, and more.

It looks like the developers are determined to continue updating the package, since everything on the site seems to have been updated within the last couple of months. They are currently on version 3.13.

This is a great little package and I intend to learn more about it and update you as I can.

The code that I used to produce the charts can be found at https://github.com/gregwa1953/FCM-171 .

```python
import plotext as plt
data=[17150,17187,17187,17187,17187,17337,17351,17358,17446,17447,17447,17447,17600,17619,17648,17694,17760,17760,17760,17876,17930,17979,18017,18061,18061,18061,18199,18222,18236,18269,18354,18354,18354,18436,18463,18510,18535,18561,18561,18561,18606,18619,18634,18666,18683,18683,18683,18731,18749,18775,18798,18816,18816,18775,18828,18840,18856,18871,18864,18864,18864,18864,18880,18904,18916,18898,18898,18898,18911,18922,18938,18951,18952,18952,18952,18974,18979,18988,18993,18991,18991,18991,19022,19034,19035,19059,19059,19059,19059,19091,19103,19126,19145,19136,19136,19136,19136,19179,19192,19217,19238,19238,19238]

plt.plot(data, label='Hays County')
plt.plotsize(100,30)
plt.title("Hays County")
plt.show()
```

This dataset has 103 data points in it. Here's what the plot looks like:
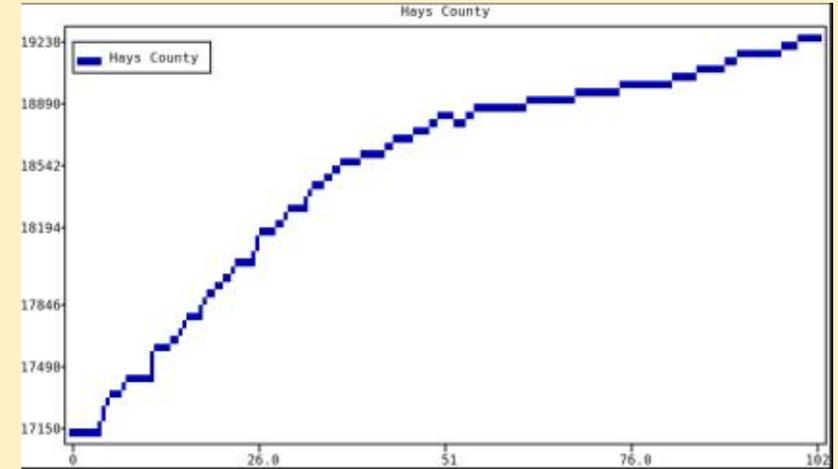
Until next time, as always; stay safe, healthy, positive and creative!

Greg Walters is a retired programmer living in Central Texas, USA. He has been a programmer since 1972 and in his spare time, he is an author, amateur photographer, luthier, fair musician and a pretty darn good cook. He still is the owner of RainyDaySolutions a consulting company and he spends most of his time writing articles for FCM and tutorials. His website is www.thedesignatedgeek.xyz .

When it comes to typesetting, LaTeX is the undisputed champ. Boika, Boika... It outweighs its opponents in the sheer number of options alone. At first, LaTeX is a bit confusing and difficult to use. Not to scare you, but I suggest looking here: http://tug.ctan.org/info/symbols/comprehensive/symbols-a4.pdf

So by this point you have an understanding of inserting pictures into your LaTeX document, but let's say you want to add some data. I am unashamedly stealing this from Wikipedia as I just needed some values (and because if I steal it anywhere else, people get upset about copyright nonsense).

So If I wanted to have a table with the maximum temperatures for British Columbia, how would you think that we would do it? With a subsection of Tables! This can be sales figures or whatever you want, we are just going to use temperatures as a neutral example. I hope you have your practise document from the last issue.

Delete everything between \begin{document} and \end{document} so we can start fresh.

This is what we will be covering if you are interested in the long explanation: https://latex.wikia.org/wiki/Tabular_(LaTeX_environment)

Our code (or skeleton, per-se) will be, if you deleted last month's file:

```
\documentclass[a4paper,
twocolumn]{article}
\usepackage{graphicx}
\usepackage{subcaption}
\graphicspath{dir-list}
\usepackage{lipsum}

\begin{document}

\end{document}
```

```
\begin{figure}
    \begin{subfigure}
        \includegraphics{Figures/you}
        \caption{It's a U!}
        \label{fig:you}...
    \end{subfigure}
\end{figure}
```

If you looked at the Wiki, you would have seen a bunch of "c's" following the word 'tabular'. Now what is all this about? Well, the 'c' is shorthand for centre, so it is the alignment of your columns. From this you should be able to figure out that 'l' will be left alignment, and 'r', right alignment. The pipe symbol is the line, separating the columns. Starting to make sense? I thought so! We generally do have smart readers.

I am going with four columns, as our stolen data has four months listed. We can do that first and see what it looks like. You can name your columns and use the ampersand "&" to separate them.

Monkey see, monkey do quickly and I will explain.

```
\subsection{Tables}

\begin{table}

\begin{tabular}{||c|c|c|c||}
    Jan & Apr & Jul & Oct

\end{tabular}

\end{table}
```

I inserted two pipe symbols on either side of my columns, to get double lines. Pipe symbol = line, if that was not clear. It looks very ugly, but it illustrates something you needed to learn. Tabular is included in LaTeX and does not

Average daily maximum and minimum temperatures for selected cities in British Columbia[29]

Municipality	January Max	January Min	April Max	April Min	July Max	July Min	October Max	October Min
Prince Rupert	5.6 °C (42.1 °F)	−0.8 °C (30.6 °F)	10.2 °C (50.4 °F)	2.5 °C (36.5 °F)	16.2 °C (61.2 °F)	10.5 °C (50.9 °F)	11.1 °C (52.0 °F)	4.9 °C (40.8 °F)
Tofino	8.3 °C (46.9 °F)	2.3 °C (36.1 °F)	11.9 °C (53.4 °F)	4.0 °C (39.2 °F)	18.9 °C (66.0 °F)	10.5 °C (50.9 °F)	13.6 °C (56.5 °F)	6.3 °C (43.3 °F)
Nanaimo	6.9 °C (44.4 °F)	0.1 °C (32.2 °F)	14.1 °C (57.4 °F)	3.9 °C (39.0 °F)	23.9 °C (75.0 °F)	12.3 °C (54.1 °F)	14.6 °C (58.3 °F)	5.2 °C (41.4 °F)
Victoria	7.6 °C (45.7 °F)	1.5 °C (34.7 °F)	13.6 °C (56.5 °F)	4.3 °C (39.7 °F)	22.4 °C (72.3 °F)	11.3 °C (52.3 °F)	14.2 °C (57.6 °F)	5.7 °C (42.3 °F)

require you to import any packages. We use {Tables} to make our tabular columns float. That way we can align our table with text. Please, at any time, feel free to play with the options you have learned so far. Make more columns, change the amount of lines, whatever makes you happy. If you make a mistake, see if you can rectify it before looking back here. Feel free to add some ipsum text, to see what it looks like.

We have columns, but how do we get rows? In LaTeX, rows are called horizontal lines, so you will need to pre-fix and post-fix each data line with 'hline'. The other thing you need to know is that two slashes "\\" end a line.

If you looked at the example on the Wiki, did you figure it out? Don't worry, me neither. Monkey see, monkey do time. I will do the first and you see if you can do the rest from our stolen table.

```
\begin{tabular}{||c|c|c|c||}
    \hline
    Jan & Apr & Jul & Oct \\
\hline
    5.6C & 10.2C & 16.2C &
11.1C \\ \hline

\end{tabular}
```

I added the Celsius symbol so you can see that it is not numeric only. Go ahead and make a table. Admit it, it was easier than you thought.

Homework: add a caption to the table explaining what we are looking at.

This is the basics of tables, we can go deeper, but we have a lot to cover, and this series is aimed at newbies who have no idea what LaTeX is. They are kept short on purpose, as LaTeX can overwhelm those who have no idea how complex it can be.

Reference for you: https://riptutorial.com/latex/example/

If you managed to keep up, well done! See you next issue with more cryptic codes.

There are more graphics options available from the tikz package. That is another kettle of fish completely. Tikz can net you a 12-part series on its own. We will maybe touch on it at the end of this series.

Did I make a mistake? Do you have questions? E-mail us at: misc@fullirclemagazine.com
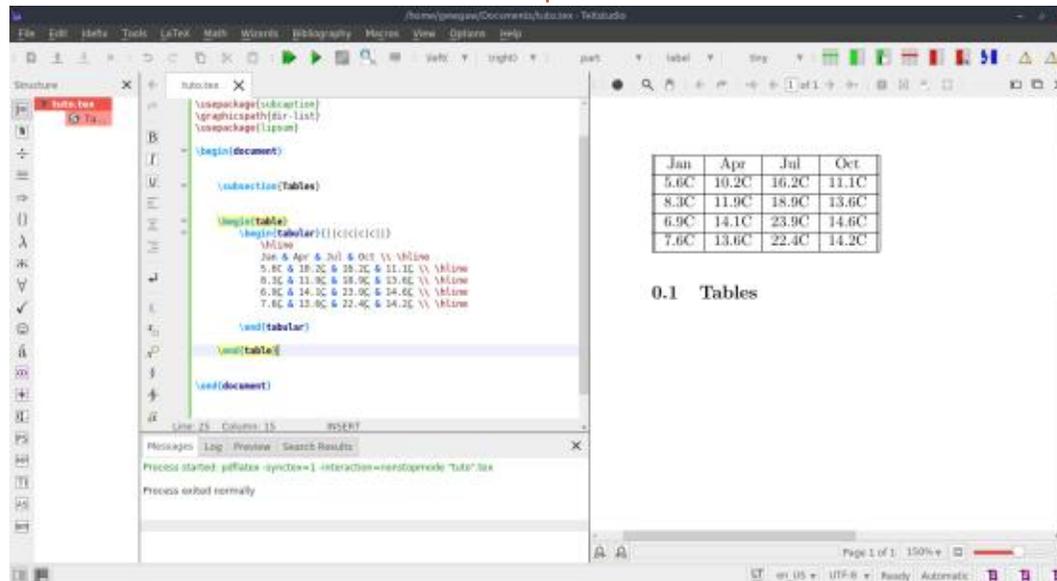
Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

When I first started using Ubuntu over 10 years ago, it was because of its ability to boot from a Live CD, thinking that this facility would be useful if I had a problem with viruses that prevented normal booting of my Windows system. I quickly realized the power and potential of Ubuntu and, although I occasionally use Windows for certain specific purposes (Income Tax, for example), Ubuntu is my working system. Many years ago, I tried, with some success, to install a full Ubuntu system on a USB stick but it was very slow, and using a Live USB was actually preferable.

I was still intrigued by the idea of having a totally portable version of my home system that I could plug into any computer. There is a substantial amount of information on this topic on the Internet. but it is somewhat fragmented and often very specific to one situation or another. Typically, it is aimed at the expert user rather than the amateur (like myself). After much research and experimentation, I have distilled what I have learned into what I believe to be a straightforward technique for creating a portable hybrid clone. I decided to try it out on a 128GB high speed USB stick I purchased on sale from Amazon (£12 or $20 CAN). With claimed read/write speeds of 300+ MB/sec, I thought that it would make a great way to run Ubuntu on my new, but entry level, Win 10 laptop which has only 64GB eMMC storage and consequently is not suitable for a dual-boot setup. Although a bit skeptical of the advertised speeds, I was pleasantly surprised to find that the system booted from the USB that I created in under 40 seconds, including login!

This article provides a step-by-step discussion of how I accomplished my goal of creating a clone of my home system (Ubuntu Mate 20.04 LTS) that will boot on either a UEFI based computer or an older machine that requires Legacy/BIOS boot. I have tried to include just enough background information to allow for an understanding of the process without getting bogged down in technical detail. I should also say that, in Linux Land, there are always many ways to do things and I make no claim that this method is the best; however, it does make use of very common Ubuntu tools and minimizes the use of the terminal. Hopefully this will help make it accessible to a wide range of readers.

It is not possible to anticipate the particulars of every possible system, so I have made a few assumptions:
• The system to be cloned is a typical installation with everything on a single partition and no swap partition (if you do have a separate / home partition, and/or /swap partition, use the technique described below to clone them also. Boot is not needed as it is re-created later).
• The removable media that you are using can be an external HDD, SSD, or flash drive, but the sector size must be 512 bytes and your internal hard drive must also have a logical sector size of 512 bytes (copying partitions across different sector sizes causes problems).
• Your home system can use either UEFI or Legacy/BIOS boot, the clone will boot in either mode.
• The destination drive must be big enough to accommodate the Ubuntu partition that we are going to clone – the whole partition, not just the used space! (Resizing your Ubuntu partition to make it smaller with gparted is quite safe – the Ubuntu installer does this when you set up a dual boot installation – BUT there is always potential for data loss so don't do it unless you have backed up anything you cannot afford to lose).
• Many of these steps require Administrator privileges, so make sure you know how to use this facility on your system as it varies a bit across Ubuntu flavours.

GETTING INFORMATION ABOUT YOUR SYSTEM

Boot your computer normally and open a terminal. Type df and look under the "Mounted On" column to find the '/' symbol. On my machine, this is opposite /dev/ sda2, and this tells me that the

system presently running is on partition #2 (bottom left). Make a note of this number from your system, the device name might change after rebooting but the partition number won't, so it is important to remember which partition you are going to copy.

CLONING YOUR SYSTEM

Now reboot your computer from a Live USB, you can use the install media for your system or any other Ubuntu flavour. Booting from a live USB is important as you are going to be working on the system partitions and it is much safer if nothing is mounted. This way we can be sure that your home system in particular is unmounted, and not damaged by anything that you do.

• Plug in your installation media to a vacant USB port.
• Open the program gparted (included as part of the live USB).

This is my "testing" laptop, and you can see that I have a 250GB drive designated as /dev/sda with 4 partitions (bottom right). I have arranged my system with 2 partitions for storage of personal files (/dev/sda3 and /dev/sda4), which means I can have my Ubuntu system partition (/dev/sda2, mounted at /) at just 30GB. The first partition is unused and is a leftover from a previous installation. If you click the pulldown menu at the top-right, you will get the partition layout for the other drives on the system, in this case the portable USB (mine is 128GB) that you will use to clone your system.
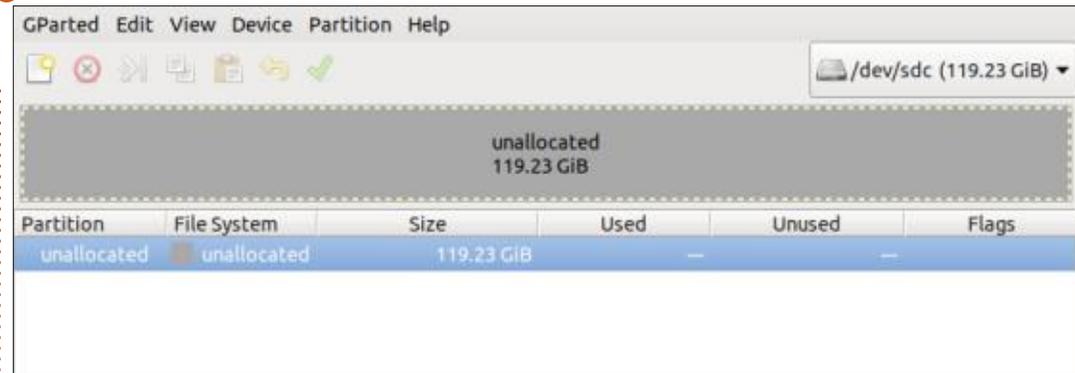
If your USB drive contains any data it will be completely destroyed by the following process so make sure you don't need anything that is left on the USB drive. The USB may show a partition but it doesn't matter, just select Device   Create Partition and make sure you choose "gpt" as the partitioning system so that it is compatible with UEFI booting. Then your drive will look like the image with unallocated space for the entire drive (top right).

Now select Partition > New to create the first of 3 partitions.

The "Free space following" will vary depending on the size of your drive, but the other values should be as shown here (below). Click "Add" and then click on the green checkmark at the top to apply this change.

contents ^
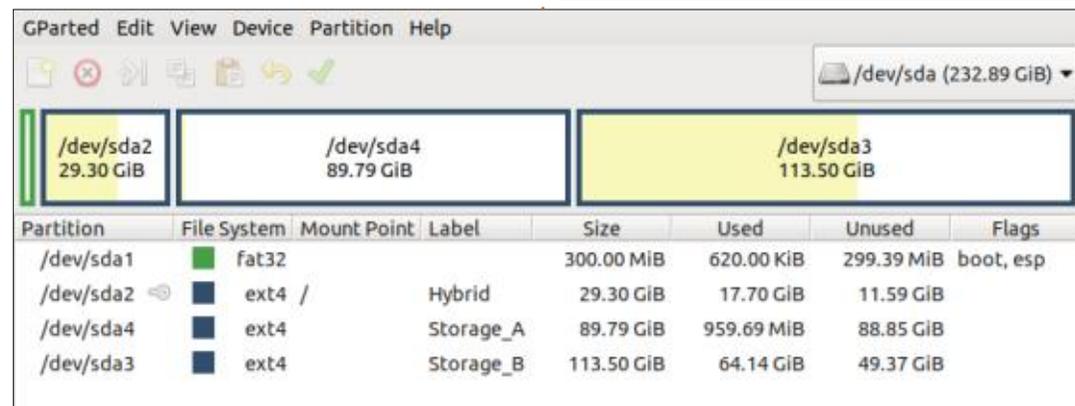
Next select Partition > New to create the second partition and adjust the values as shown:
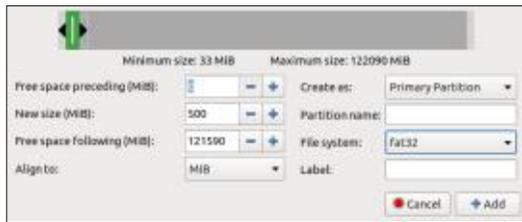
Once again, select Add and then click on the green checkmark to apply this change.

Finally, select Partition > New to create the last partition, this one for the Ubuntu system itself:

Click Add and select the green checkmark to apply this change.
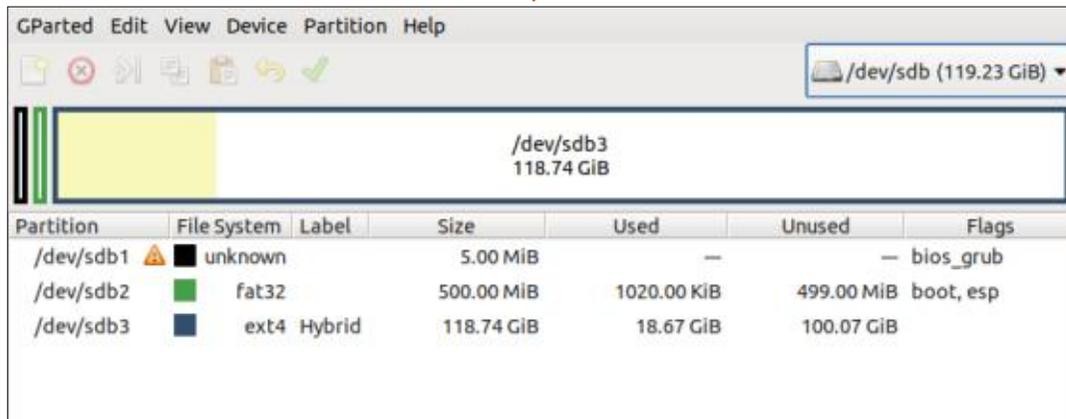
The last thing you need to do is to set the appropriate flags, so right-click on partition #1, select Manage flags, and choose bios_grub as the flag, click Close, and this change will be applied automatically. Repeat for partition #2, but this time select boot, esp as the flags. Your USB drive should now be partitioned and ready to receive its own operating system.

Each of these partitions serves an important function:
• Partition #1 provides a protected place for GRUB files used in Legacy/BIOS booting so there is no conflict with the space used by the GPT partition table.
• Partition #2 provides a place for the UEFI GRUB bootloader files to go.
• Partition #3 contains your Ubuntu filesystem.

Now go back to select /dev/sda and right-click on your Ubuntu partition, /dev/sda2 in my case, and select Unmount (if it is not mounted this choice will be greyed out.) Once it is unmounted, right-click and select Copy. Switch to the USB drive, make sure Partition #3 is unmounted, right-click on it, and select Paste. Depending on the size of the partition you are copying, this could take an hour or more, for this size of partition it was about 15 minutes. This procedure makes an exact and complete copy of your Ubuntu partition, including the label (Hybrid in this case) and the UUID (Universal Unique ID) (bottom left). Once this copying step is complete, right-click on partition #3, and give the partition a different label, for example, Ubuntu_USB, just to make it easier to identify; and, click the green arrow to apply the change. If you open a terminal and type sudo blkid, you will note that both /dev/sda2 (the hard drive Ubuntu partition) and /dev/sdb3 (the partition on the USB – the different letter is due to a reboot between taking screenshots for this article) have exactly the same UUID (shown bottom right). This is a problem for GRUB as it typically will boot the first partition it finds with the UUID in its configuration file, generally the hard drive, resulting in a failure to boot the USB version of Ubuntu.
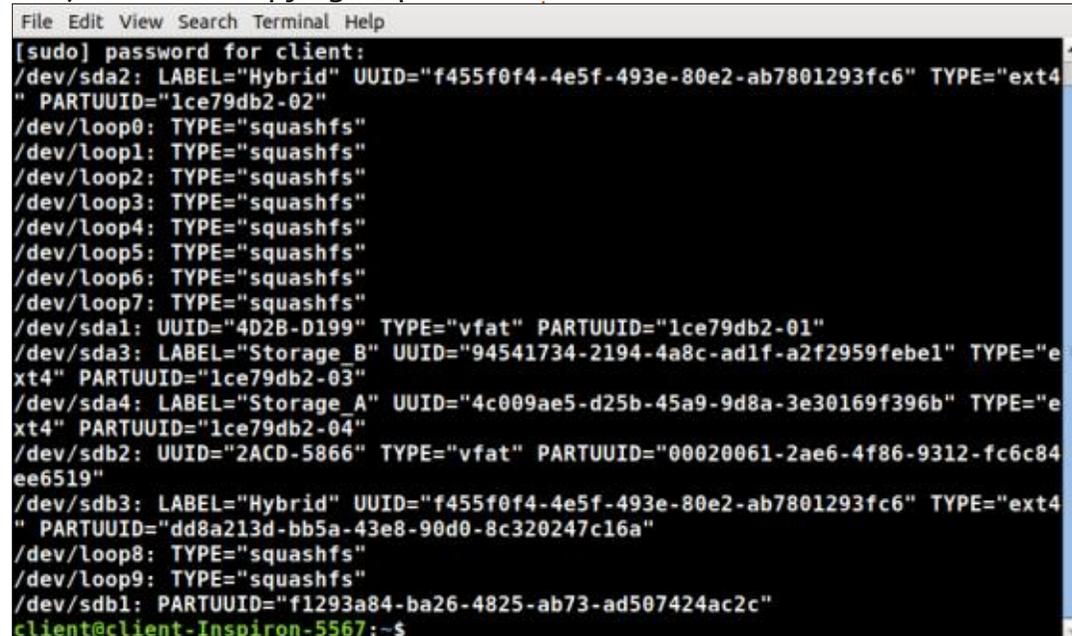
Run gparted again, select the USB drive, right-click on the Ubuntu partition (/dev/sdb3), and select New UUID. Click on the green checkmark to apply this change.

Open a terminal and type sudo blkid again and you will see that the UUID for partition /dev/sda2 is unchanged BUT /dev/sdb3 now has a different UUID (bottom left).

This change needs to be added to the fstab file. Open the file manager and double-click on Ubuntu_USB, then Edit as Administrator the file /etc/fstab. Copy and paste the UUID for /dev/sdb3 into this file and delete the previous UUID entry.

You will also need to add an entry (line 14) so that the UEFI partition gets mounted (top right).

(Note: All of these UUIDs will be different on your system so adjust as needed.) Save the file and exit.

This change in UUID also needs to be reflected in the /boot/grub/grub.cfg file, or you will be unable to boot into the USB-based system. Edit this file as an administrator and find the old UUID and highlight it (bottom right). Then select Search, and this highlighted value will be in the "Search for" box, paste in the new UUID from the terminal into the "Replace with" box, and select
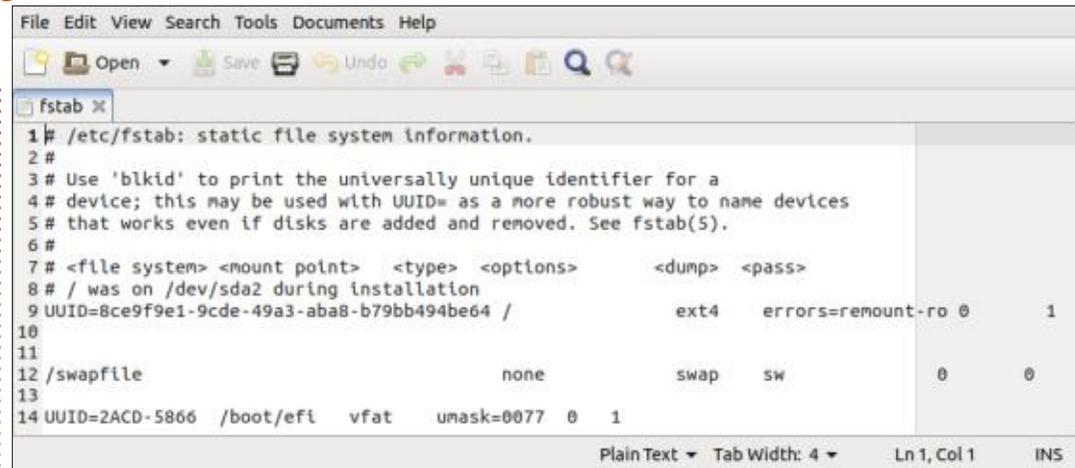
replace all. Save this file and exit. Finally, use the file manager to examine the contents of the /boot folder on Ubuntu_USB. If there is a folder called efi already, make no changes, otherwise right-click to create a new folder, and then rename it as efi. This ensures that

there is a mount point of /boot/efi for partition #2 as required.

INSTALLING BOOTLOADERS

We now have a complete Ubuntu installation on the USB but installing GRUB on removable

media has known problems – difficulties that can be avoided if we install the GRUB bootloader files from a running system. And, there is a simple way to boot into the USB system so that we can install GRUB.

Since we are now finished with the live USB, shut down the system, remove the live USB and reboot to your base system. Open a terminal and run sudo update-grub. When you do this, GRUB should find the Ubuntu installation on /dev/sdb3 (bottom left), and add it to the boot menu. It will be bootable from the GRUB menu (because we adjusted the grub.cfg file to point to the correct UUID) despite having

no bootloader of its own. If update-grub does not find Ubuntu on /dev/sdb3, you need to edit the etc/default/grub file as administrator and add or edit the following line: GRUB_DISABLE_OS_PROBER="false" and try updating GRUB again. Assuming it works this time, reboot your system and check that the new GRUB menu includes an entry for the Ubuntu system on /dev/sdb3. Use the down-arrow to select it and the system should boot into the Ubuntu filesystem on the USB drive. You should open a terminal and type df just to check that /dev/sdb3 is mounted as root to ensure that you really have booted from the USB rather than your hard drive. At this point, I suggest that
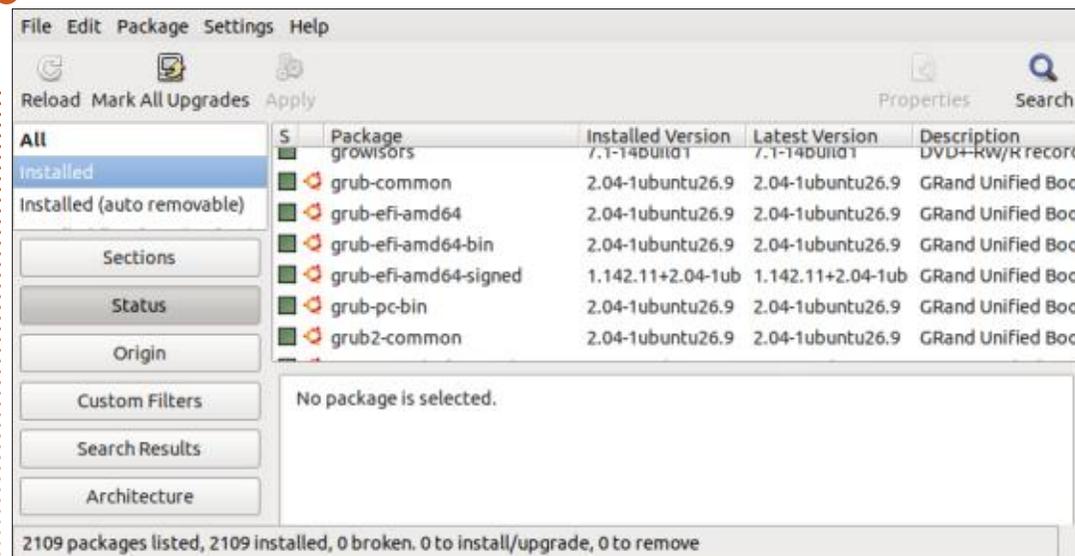
you right-click on the Desktop and change the background. Later on, this is an easy way for you to know you booted from the USB rather than the system on the hard drive. Now we are ready to install the appropriate bootloader files and check that the new portable drive works as it should.

The first thing to check before proceeding is that all of the appropriate grub packages are installed on your system. To do this open Synaptic, select Status Installed, and scroll down to grub entries. There are 6 things to look for (top right). Depending on how your system was configured these may or may not be installed but they all need to be there to proceed.

If any are missing, select All, and scroll down to the grub packages and install it/them. If you do one at a time, others (that are missing) may get installed automatically. If Synaptic is not on your system, install it with sudo apt-get install synaptic; however, if you have another preferred method for checking installed packages and installing any missing ones, you can use that instead.

For legacy/BIOS systems, open a terminal and type: sudo grub-install –target=i386-pc /dev/sdb, since the device name given in our df command was /dev/sdb3 and this tells us the drive name is /dev/sdb (next page, top right).
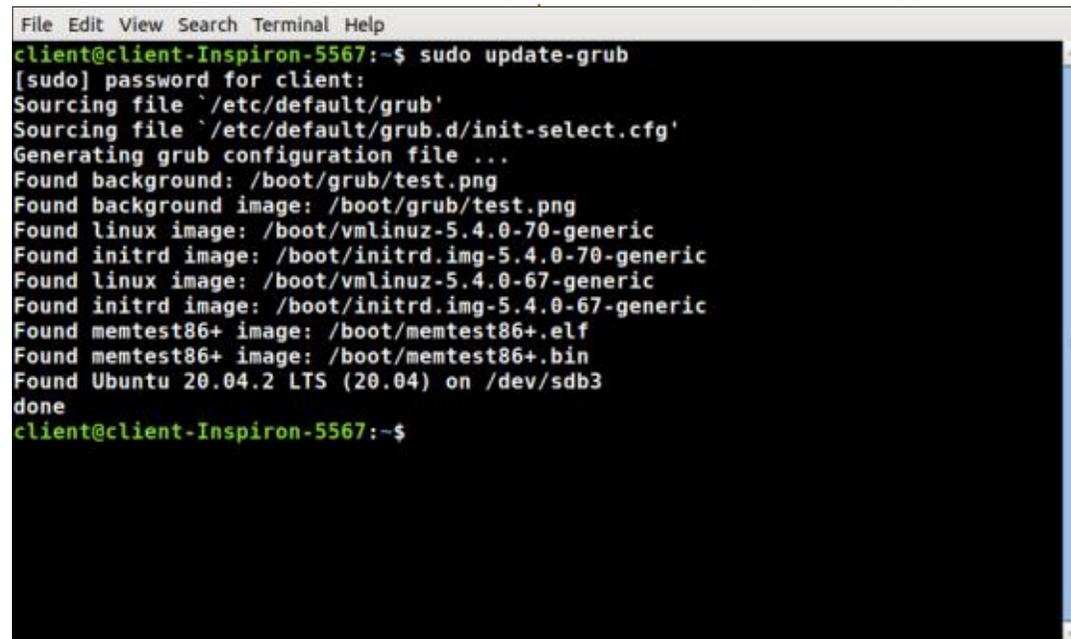
The USB can now be booted in Legacy/BIOS mode.

For UEFI systems, open a terminal and type sudo apt-get install shim-signed. (If you are told this is already installed, then use sudo apt-get install --reinstall shim-signed to repopulate the efi folder). You will be asked to confirm where you want to install the bootloader files, and this is /dev/sdb2 for this system (the UEFI partition we created at the start). On your system, the drive letter may be different so use what you found from the df command. If you are presented with a graphical screen with several choices, use the space-bar to toggle the asterisk indicating your choice, the arrow keys to navigate; and the enter key when you are ready to proceed.

The advantage to this particular method is it actually installs a series of programs that permit booting with secure boot enabled. If you Open as Administrator the folder /boot/efi, and then double-click on EFI, you will see 2 folders, BOOT and Ubuntu. These folders contain the files needed to boot Ubuntu, and also a "fallback" option in case something doesn't work properly. The final step is to type sudo update-grub which will ensure that all the pointers are correct in the grub.cfg file. You will notice that the update "finds" a Ubuntu system on the hard drive (/dev/sda2 in my case, yours will be different), and this will be reflected in the grub menu with a boot entry. You can just ignore it, but, if you are offended by it, edit the file /etc/default/grub (as administrator), and add (or edit) the line which invokes the os-prober to search for other operating systems to read: GRUB_DISABLE_OS_PROBER="true". This will prevent other operating systems being detected and added to the grub menu. You will need to run sudo update-grub after you do this so that it takes effect – this might result in the boot menu not being displayed when you boot since there is no other system to choose; but it will still wait the specified timeout behind the scenes.

Your hybrid clone should now work. Reboot your system and press the F12 key (or whatever key your computer uses) to bring up the boot menu. Select the USB device and you should boot into Ubuntu with the new background screen (this way you know you aren't booting from the hard drive).
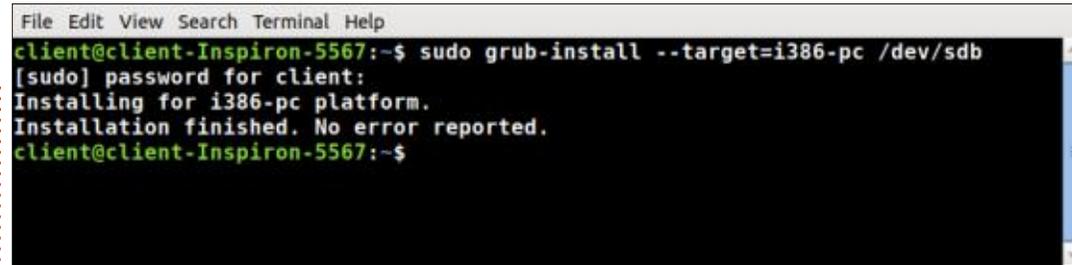
```
File  Edit  View  Search  Terminal  Help
client@client-Inspiron-5567:~$ sudo grub-install --target=i386-pc /dev/sdb
[sudo] password for client:
Installing for i386-pc platform.
Installation finished. No error reported.
client@client-Inspiron-5567:~$
```

I have successfully booted this USB flash drive on 5 laptops/netbooks of varying age and type (Acer, Asus and Dell) as well as both Legacy/BIOS and UEFI boot – with and without Secure Boot enabled. It is impossible to guarantee that it will perform correctly on every system; but, hopefully it works for you!
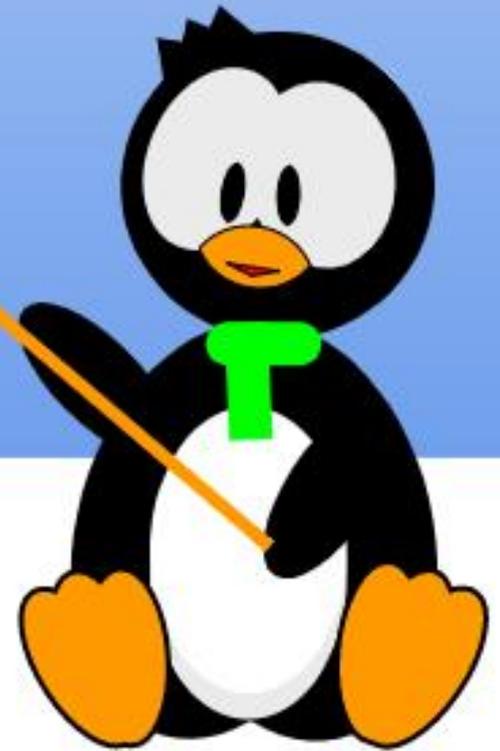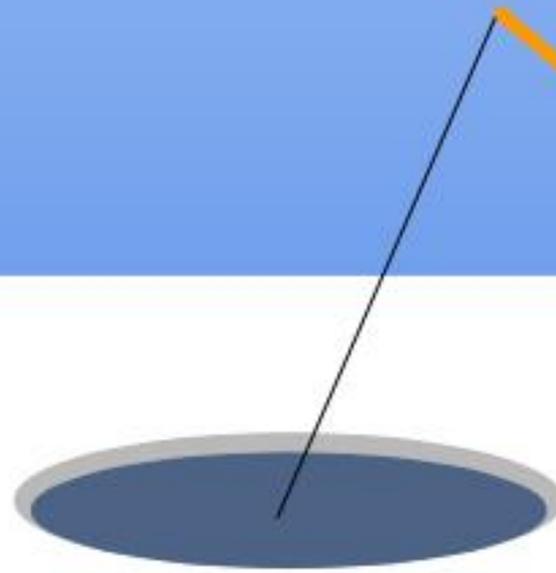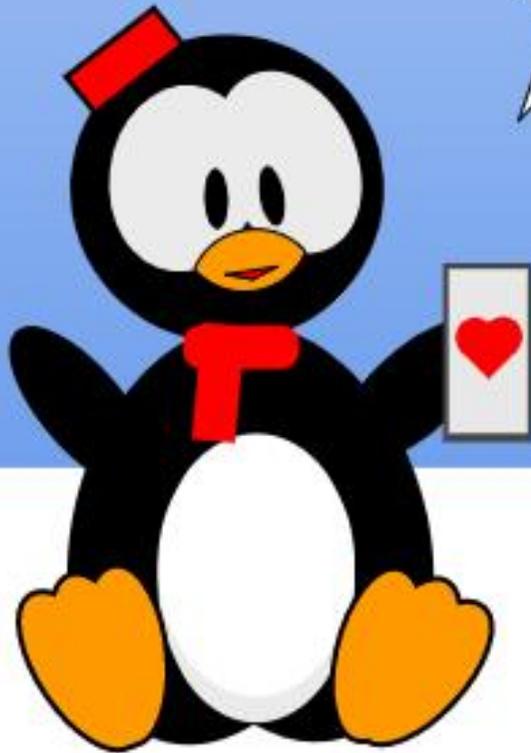
This month, I'll be concluding my look at the Trace Bitmap dialog by examining the new tracing modes that have been introduced with Inkscape v1.0, including one that has been much requested on the forum over the years.

AUTOTRACE

Inkscape has long offered tracing of bitmaps using the Potrace library. While this does a fine job of vectorising raster graphics, it's not the only open source tool that performs this trick. Autotrace is a command-line tool of similar vintage, which has now been
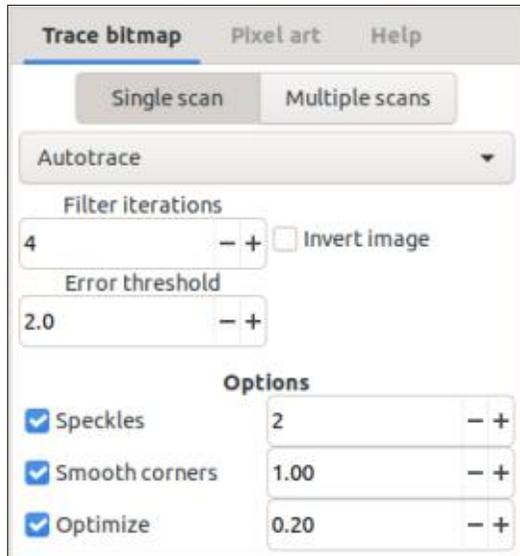
integrated into the Trace Bitmap dialog. It's available in the mode selection pop-up, in both single-scan and multiple-scan varieties. The latter is labelled as "Autotrace (slower)", suggesting that it probably shouldn't be your first choice for a multiple-scan conversion – a suggestion that I'll come back to later. For now, however, we'll start by looking at the single-scan variant.

Once selected, it offers two spinboxes and an "Invert image" checkbox. The latter, as the name suggests, inverts the image colors before tracing, making it easier to trace a light-on-dark design without having to manually process it first. The effect of the two spinboxes is somewhat harder to guess at.

The scant documentation for Autotrace – basically an online copy of the man page for the command-line tool – describes the "Filter iterations" option thus: Smooth the curve the specified number of times prior to fitting. Without knowing what is meant by "fitting" in this algorithm, it's hard to infer

what this actually does. Trial-and-error suggests that it reduces the fidelity of the trace somewhat, smoothing out the generated vectors and tending to result in a trace with fewer nodes. Here's what the most extreme values look like when applied to dear old Frankie. The original bitmap is on the left, with clipped versions of the traced head to the right. For the top trace, the filter iterations was set to 1; for the bottom it was set to 20.
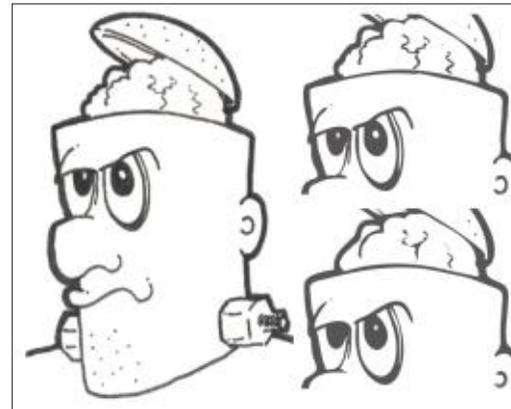
Note that the higher iterations have reduced or wiped out the whites of the eyes, while the veins on the brain are a mere shadow of the original. The shape of the ear and of the brain's outline are also

significantly smoother. The total node count for the traced head is 485 for the top image and 410 for the bottom one.

The second spinbox, for setting the "Error threshold", is described like this in the man page: Subdivide fitted curves that are offset by a number of pixels exceeding the specified real number. This one I can at least have a guess at. The process of tracing a bitmap consists of generating paths that approximate the shape of the original pixels. The paths will most likely be a close match in some areas, but not as accurate in others. This parameter allows you to set the threshold before which a path segment will be considered too far away, and will be subdivided into two paths to make it easier to adjust them to fit. Setting a small value here allows only slight deviation from the pixel positions, at the expense of a lot
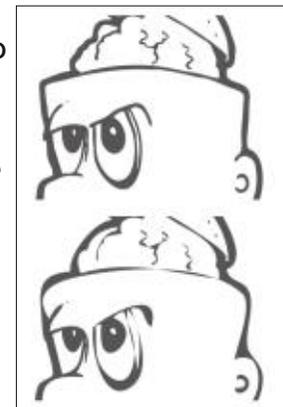
more subdivisions and, therefore, more nodes in the result. Let's take a look at how Frankie fares with values of 1.0 (top) and 10.0 (bottom).

It's pretty clear that the higher value results in a trace that is so smoothed out as to lose many of the original shapes completely. The top image, where the paths were much more heavily subdivided, consists of 587 nodes; the bottom one has only 327.

As is often my advice in this column, I suggest most users should at least start with the default values for both spinboxes, and start tweaking them only if you need to improve the fidelity of the trace, or want to take the counter-approach of reducing the number of nodes. Even in the latter case, I would probably be more inclined to trace with the defaults and then use Path > Simplify afterwards.

Perhaps the biggest question is how the Autotrace results compare with the Potrace equivalents. Here's another pair of traced Frankies created using the default settings: The Potrace-based "Brightness cutoff" at the top; the

Autotrace version at the bottom. Again, the full head on the left is the original raster image.
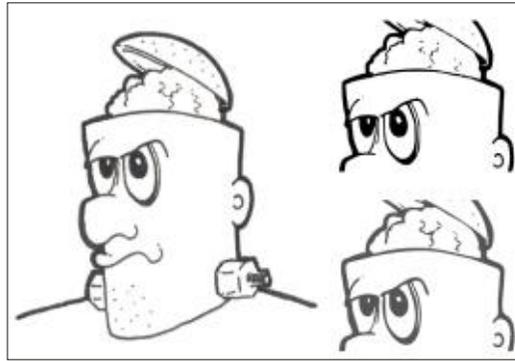
The first thing to note is that the Autotrace version has maintained the grey color of the original image – though that's such a trivial thing to change that it shouldn't be used as a reason to select one over the other. The Potrace result is a lot crisper, with the paths more accurately maintaining the sharp corners of the head, and thinner lines of the eyebrows. This accuracy is reflected in the node count: 1090 for Potrace but only 440 for Autotrace.

But it's not that clearcut. The extremely thin lines on the brain are actually better preserved by the Autotrace algorithm. On the whole, I think the old Potrace code works best, at least in this case. But I also wouldn't rule out creating a hybrid result by using node editing or

Boolean operations to paste together the best parts from each result.

AUTOTRACE (MULTIPLE SCANS)

What about using the "Autotrace (slower)" mode for scanning color images? My advice is to avoid it completely and stick to the Potrace-based modes. I tried scanning the same images that I used for part 19 of this series: the Full Circle Magazine logo, and a Wikimedia Commons copy of "La Giaconda" (The Mona Lisa). In both cases I used the default settings. The logo, which takes less than a second to trace with Potrace, took several minutes to complete. With such an amount of effort involved you might expect something impressive, but this is what the result looks like (original bitmap on the left, Autotrace in the middle, Potrace on the right):

To you and I it may appear as though Autotrace spent several minutes producing a salmon-colored circle. But no: what you're actually looking at is a group of 4180 objects! For comparison the Potrace version contains 8 objects – one for each color set via the "Scans" spinbox.

Switching to the outline view does suggest that the shapes have been traced, and are hidden somewhere in the salmon fillet before us, but the thicker outlines definitely hint at complex paths compared with the simplicity of the Potrace version.
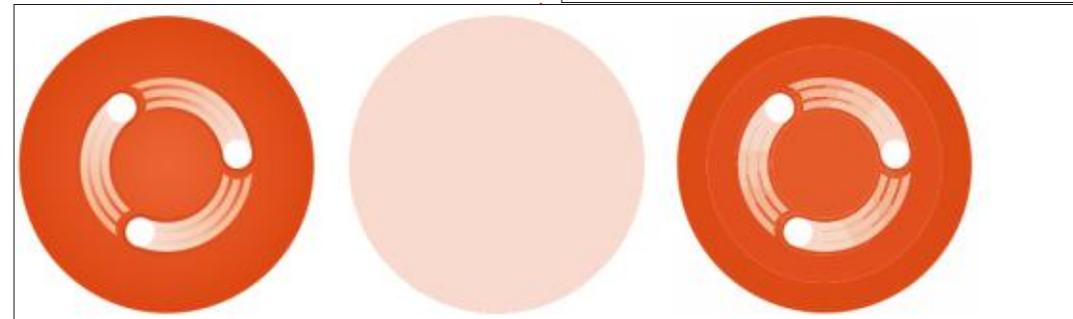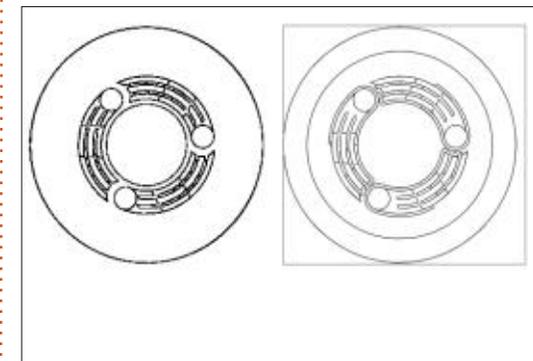
And what of La Giaconda? After many minutes of processing my memory, swap and CPU were all maxed out, then Inkscape disappeared off my screen entirely. There was no appearance of the usual crash dialog I see when it dies, leading me to suspect that its demise was perhaps the fault of the Linux kernel killing it due to lack of available resources.

Even with a "successful" trace, the sheer number of objects created is practically unmanageable. There may, perhaps, be some image types for which this mode offers an advantage, but I would try it only if the Potrace methods aren't yielding acceptable results – and make sure to save your file first!

CENTER LINE TRACE

If Autotrace offers little or no improvement over Potrace, and in some cases is far too resource hungry, why bother adding it to Inkscape at all? The reason is that it offers one type of frequently requested tracing mode that Potrace does not – center line tracing. In fact the menu entries described previously are there only as a side-effect of including this mode. After all, if you're adding the library anyway, why not also expose the standard tracing mode as well, to give your users more options.

Center line tracing is really applicable to only line art in which the shapes are made up of individual pen or pencil strokes. Using other tracing modes, each stroke is converted into a closed, filled path that reflects the thickness and shape of the original artwork. With this new mode, however, the tracing algorithm attempts to determine a single path that traces out a line following the middle of the original stroke.

For the most simple real-world example, consider a single pen stroke on paper, scanned and imported into Inkscape. The top line in this image is the original scanned raster graphic. The second shows the result of a normal trace – note that the bulbous ends of the line are reproduced in this mode. The third line is the result of a center line trace – no thickening of the line at the ends of this version.
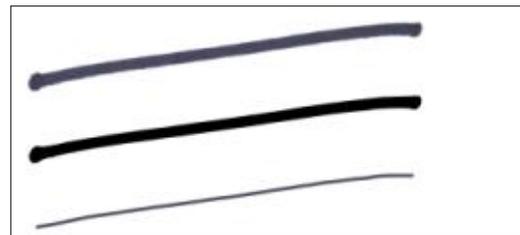
The real difference becomes clear when we take a closer look at the nodes used to make up the two traced paths. The first is a closed, filled path, so you can see that the nodes make up the outer shape of the stroke. The center line trace, on the other hand, results in an open path made up of a simple line of nodes: any suggestion of line thickness is purely down to the value set for the stroke width.
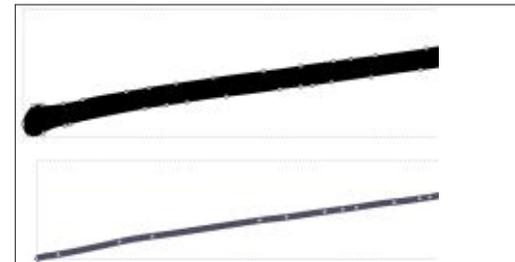
What happens when you try this mode with a more complex example? How about a few handwritten letters?

As you can see, the trace doesn't really reflect the shapes and writing style of the original scanned image. The fairly straight left leg of the A becomes kinked in the middle, as the algorithm struggles to work out where the center line actually is. The sharply angled line where the top and bottom bowls of the B meet is lost entirely, and replaced with a horizontal crossbar.

You may also have realised that such shapes can't be made up of a single path segment. In this instance we've ended up with a single complex path consisting of all the different segments that make up the letters combined into one object. Path > Break Apart allows us to reduce the complex path down to its constituent parts, which we can then give different colors to demonstrate the paths that the algorithm settled on.
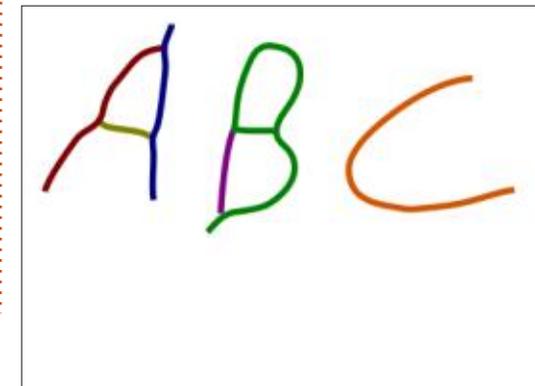
The A, not unreasonably, is made up of three separate paths. But the complexity of the B is captured in only two paths: one complex curving line that encompasses most of the shape of the letter, and a small straight segment to fill in the remaining gap. Wrangling such shapes into something more befitting the original outlines could quickly become tedious on larger projects.

You may think I'm being unfair on the algorithm here. My scanned text was from a thick Sharpie, rather than the thin strokes of a pencil or ballpoint pen. But based on my testing, you'll likely face similar issues, even when starting with thinner lines in the source material.

That's not to say that the new mode is useless or unwelcome. For many images it will prove to be far more effective than the existing tracing methods, especially if you're interested in only the core shapes of the elements rather than the exact details of the stroke outlines. Just remember that it's working only with pixels, and has no concept of the order in which lines were laid down, or the difference

between two lines that meet at an angle compared with a single line that has a sharp corner in it.

As is so often the case with the Trace Bitmap dialog, I can only recommend that you give it a try on your image, but don't expect miracles. Even if the results aren't perfect, it may save you some manual tracing time on part of your design, or at least give you a starting framework to build upon.

Next month, we'll take a look at the new "Selectors and CSS" dialog, which promises to make Inkscape a little more useful as a web development tool.

LINKS

Potrace: http://potrace.sourceforge.net

Autotrace: http://autotrace.sourceforge.net/ https://github.com/autotrace/autotrace

Autotrace man page: https://linux.die.net/man/1/autotrace

"Frankie" and other images: http://www.peppertop.com/fc/

Mark uses Inkscape to create three webcomics, 'The Greys', 'Monsters, Inked' and 'Elvie', which can all be found at http://www.peppertop.com/

Well, it's time for a return to retrogaming as a topic. Honestly, I wish I had enough time to write a separate monthly retrogaming column because there is SO MUCH that could be covered. Still, we can revisit this topic from time to time in Everyday Ubuntu, and it has been a little while since our last look at getting old DOS games running on modern Ubuntu.

Back in 1992 or 1993 or thereabouts, I was working at Dallas Semiconductor when one of my bosses asked me how I came to know so much about computers, and I told him that learning how to run games was probably the single biggest factor. You see, back in those pre-Windows 95/pre-Linux GUI days, you had to work surprisingly hard to accomplish much of anything on a computer. Many DOS games required knowing how to manage memory, assign IRQs and DMAs for hardware, edit configuration files, and more. Linux was even more challenging. I remember having to learn the rudiments of the included C language editor and compiler just to install a really basic Pac-Man style game.

Although things are MUCH easier now on both Windows and Linux, getting old DOS games to run on Linux is a satisfying and enjoyable activity, a challenge of its own, and a way to learn more about how Linux actually works, as we'll see in this month's Retrogaming Revisited.

DUNGEONS AND DRAGONS: THE PRINT GAME

In the early 1980's, my good friend Trey Seastrunk (who remains a good friend to this day), talked me into going to visit a guy who was Dungeon Mastering a D&D session. I was somewhat reluctantly roped into playing in the session myself, although I knew absolutely nothing about the game. Much to my surprise, with minimal coaching, I was able to contribute to the party's adventure successfully and decided it seemed pretty fun. Soon I was buying the first edition Advanced Dungeons and Dragons core rulebooks (Dungeon Master's Guide, Player's Handbook, and Monster Manual), and creating my own fantasy world as a backdrop, where I would create my own adventures and occasionally drop in pre-designed game scenario modules that I would purchase.

Over time, TSR, the company that published D&D, put out quite a lot of modules, from the basic adventure background location 'The Keep on the Borderlands', to the notorious player-killing meat grinder of a dungeon called 'The Tomb of Horrors'. I've always wished I could play Tomb of Horrors as a player to see how challenging it really is. I have the module and did read through it, but that was literally decades ago and I don't remember any specifics, so it would be a legitimate challenge.

And there's actually a way to do this, without having to find someone willing and able to DM, or Dungeon Master, a game session. There's even a way to do it specifically on Linux.

DUNGEONS AND DRAGONS: THE GOLD BOX SERIES

From 1988 to 1992, SSI (Strategic Simulations Inc), under license from TSR, produced a series of D&D games for DOS (actually, originally for the Commodore 64) called the Gold Box series. They literally came in recognizable gold-colored boxes, used a common game engine, and allowed you to play D&D on your computer, even with no other players around. SSI created an ongoing story with these games, set in a detailed and lively game world. Titles included Pool of Radiance, Curse of the Azure Bonds, Hillsfar, Gateway to the Savage Frontier, and quite a few others, as the series was a roaring critical and commercial success.

However, even this success had its limits, and by 1993, TSR had a condition for renewal of the license: retire the now clearly aging and obsolete Gold Box engine. SSI responded by creating a fantastic swan song – Forgotten Realms:

Unlimited Adventures.

Rather than being just an actual game, FRUA was essentially a game-creation toolkit. Users could create their own 'Gold Box' adventures and load them into FRUA. A still active online community grew up around this, and soon there were hundreds of modules free to download and play for anyone who owned Unlimited Adventures, including adaptations of many of the print modules.

Enter: GOG.COM

Fortunately for us Linux users, we don't even need to manually use DOSBox and try to find an old copy of FRUA on EBay or Craigslist.

GOG.COM has a version of FRUA remastered for Ubuntu, and it's currently on sale (at the time of this writing), as part of their Forgotten Realms Collection 2 (which includes several other Gold Box games) for the only mildly outrageous price of … two dollars and forty-nine cents. Definitely worth the cost, in my humble opinion. I wouldn't remotely balk at paying the non-sale price of $9.99; it's really great value even at full price. End of sales pitch, I don't exactly get a commission….

Installing the Game: Bash Shell Scripts

For me, I've actually owned this package for years, but have never installed it or tried it on Ubuntu. Time to remedy that omission, I think! We'll start by going to GOG's page for owned games and download the installation program to a local folder. Here, mine is in the folder $Home/Downloads/Games/ FRUA:

As you can see, the file is of type and extension sh. This means it is a bash shell script file. Bash is the normal default shell for Ubuntu Linux, meaning that when you go into the Terminal to perform command-line functions, you are actually running Bash in order to do so. There are other shells that can be used, like KornShell, but for most users, even professionals doing high-end functions in the command-line interface, Bash works just fine. For our purposes here, Bash is more than sufficient.

A Bash shell script is a file that is a sequence of commands, in a plain text format, telling Linux to perform a function or (more often) a series of functions. In this case, GOG has provided an installer for Forgotten Realms Unlimited Adventures in the form of a script. But, as it stands right now, Linux does not understand this by default. It will, if you double-click it, open the file for editing:

Fortunately, GOG has been good enough to provide instructions at
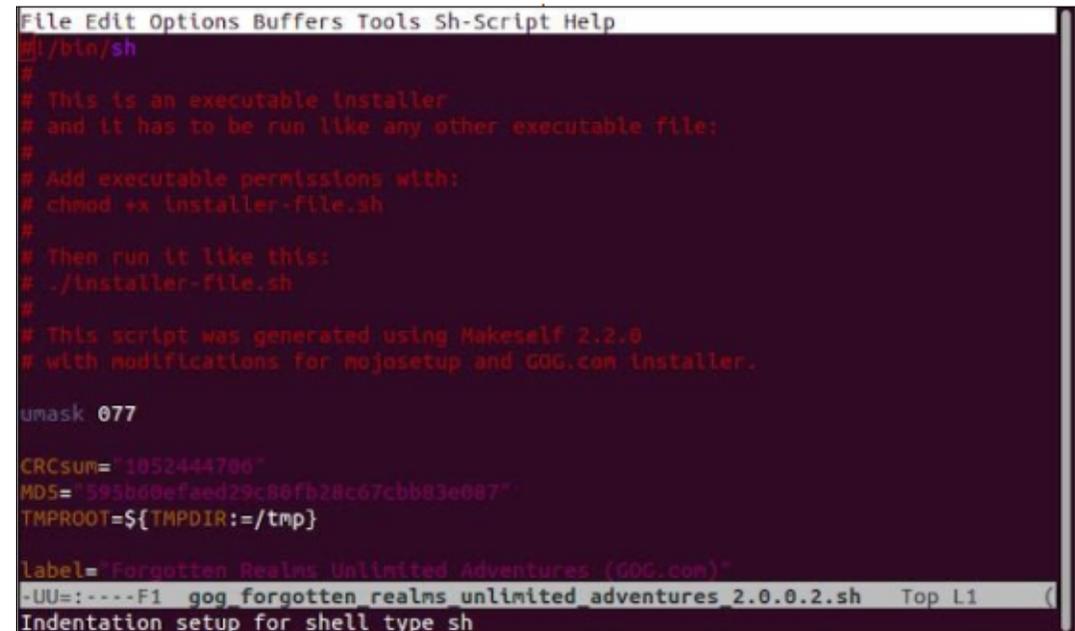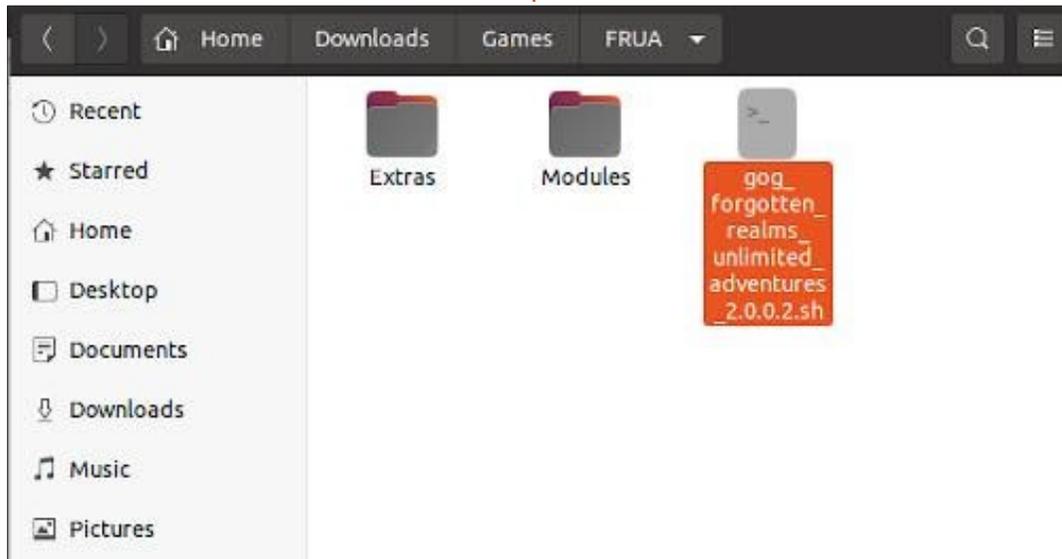
the top of the text file. Lines in a script file like this that begin with a # character are ignored when the file is run as an executable script, so those lines are used to document commands and functions within the file. They are equivalent to REM command lines in old-time BASIC computer language interpreters, and are an important way to inform others of what your intent is if you create a script file. "But, what is this chmod command the file mentions, and how do I access it?" you ask with pleading eyes and trembling voice. Not to worry, beloved reader, you actually don't need to know. CHMOD is a very powerful and useful command, to be sure, but your average Ubuntu end-user will probably never need it. Instead, we'll get to where we need to go using the File Explorer.

What we're actually trying to achieve is to let Ubuntu know that we want to execute this shell script file, not to edit or change it. We can do this without using chmod. What we need to do is to right-click the .sh file in the File Manager and go to Properties.

Go to the Permissions tab at the top and look on the line that says Execute. Click the checkbox next to 'Allow Executing File as Program':
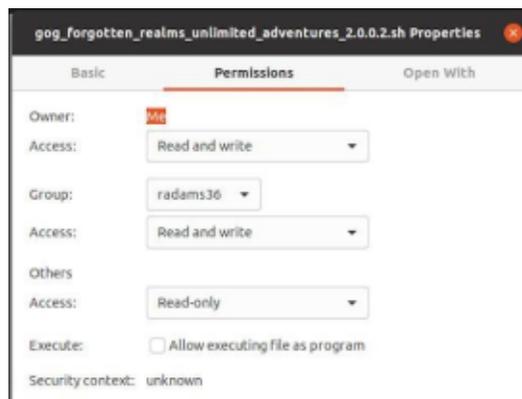
This will accomplish what is otherwise manually done by using the chmod command as the script requests. Now, you can double-click the file and the script will run as an executable file, once you confirm that you do want to run the script:
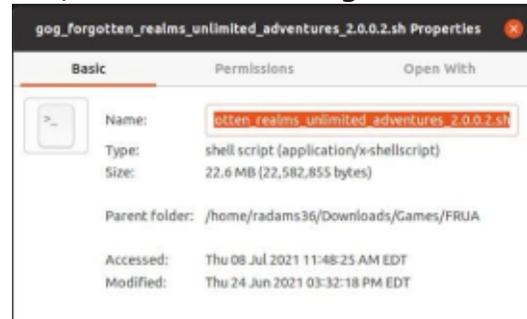
Huzzah! Now we're in the executable installer!
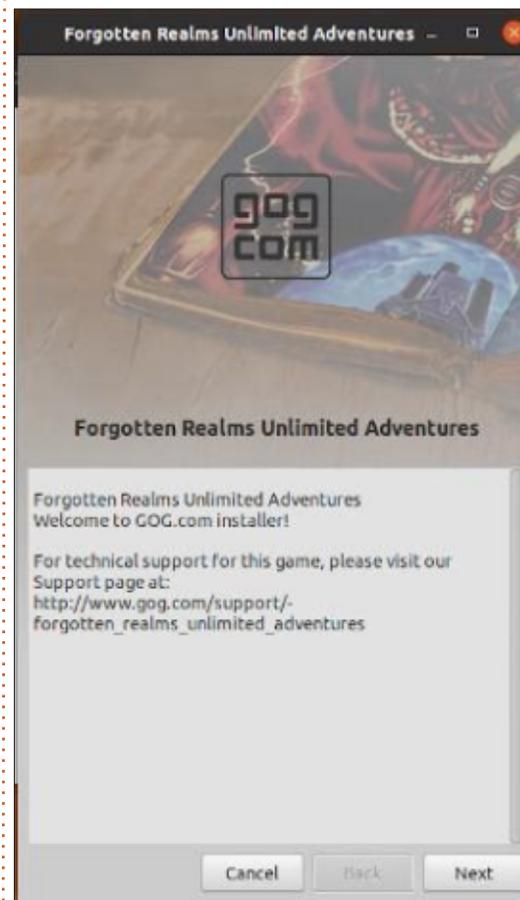
Next month: We'll complete the installation, download and unzip the required support files, and discover how to run the many downloadable modules for the game.

Richard 'Flash' Adams spent about 20 years in corporate IT. He lives in rural northwest Georgia, USA, with his adopted 'son', a cockatiel named Baby who now has a little brother, a dusky-headed conure named Skittles. Feedback welcome at: acer11kubuntu@gmail.com

This month, we will be talking about using the WS2812 Integrated Light Source on the Raspberry Pi Pico, and on the ESP32 and ESP8266 microcontrollers. The sharp-eyed among you who are good and loyal readers, might be thinking that, last month, I said we would be talking about NeoPixels. You are correct. I did. And we are.

NeoPixels is a brand name that belongs to Adafruit. It stands for the individually addressable RGB color LEDs; strips all based on the WS2812 and WS2811 LED with drivers that use a single wire protocol. So, TECHNICALLY, not all WS2812 LEDs are NeoPixel displays.

The ESP32 and ESP8266 microcontrollers both have a driver that is included in the distribution of MicroPython. For the Raspberry Pi Pico, the communications to the LEDs needs to be done via PIO programming - Raspberry Pi's assembly language. Luckily, there is a library available for the Pico to handle the communications. You can find it at: https://github.com/benevpi/pico_python_ws2812b.

This library works very well, but has been forked and updated. However, the new libraries don't work correctly from the Thonny IDE, and it's a real pain to write your code, and then try to run your file outside of the IDE. We'll use the ws2812b library for now.

SETTING UP YOUR BREADBOARD

NeoPixel type devices take a good bit of power. If you are using only a single LED, then it's not too big of a deal, but if you have 8, 24, 30 or 60 LEDs, you will need an external power source of about 5 volts DC. I say about 5 volts, since 5 volts is the absolute maximum voltage you should provide to the NeoPixel type devices. The current requirement of multiple LEDs is surprising, and if you attempt to use the USB power to drive the LEDs, it will quickly become too high for the PC to provide.

I use a 3 x AA rechargeable battery pack for my setup, with NIMH (Nickel Metal Hydride) batteries. This provides about 4.6 volts (1.2 volts each) of power when the batteries are fully charged. This is low enough to provide enough, but not too much, power.

While the connections for the RPi Pico and the ESP32/8266 are very similar, I will include one of each. In the illustrations, I will use a WS2812 8-LED stick. This stick has a fairly common set of connections that will easily work for the stick and most any flat strip, no matter the length and number of LEDs. All devices will have at least 3 input pins. Ground, +5 volt DC, and a pin

for data in. No matter if you are using an 8-LED stick or a 1 metre 60-LED weatherproof strip, the connections are basically the same.

As you can see in the image below, there are four inputs on each side of the stick. On the left, there are 2 for ground, one for data in and one for the +5 volt power. The right side has the same pins with the exception of Data Out, which is used (if needed) to add another stick.

One other thing to note is that the breadboard images show a 3 x AAA battery pack. It should actually be a 3 x AA battery pack. I couldn't find the actual image in Fritzing.

RASPBERRY PI PICO

For the Pico, we put the 3 x AA battery pack to the + and minus rails of the bread board. The data-in pin of the LED stick is connected to physical pin 21, which is GPIO 16. The Ground and +5 volt pins of the stick are also connected to the power rails of the breadboard. Finally, be sure to connect the

8 WS2812 LED Stick

Ground
Data In
+ 5 volts
Ground

Ground
Data Out
+ 5 volts
Ground

ground rail to a ground pin of the Pico. It is important that all devices share a common ground reference.
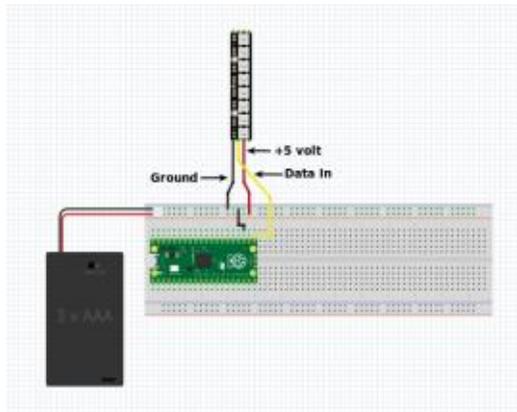
ESP32/8266

The ESP series breadboard connections are very similar to the Pico. The main difference is where the ground and data pins are. You can use any ground pin for the ground, and any GPIO pin for the data pin. In the case of this image, I decided to use pin # 27 (GPIO 16) as
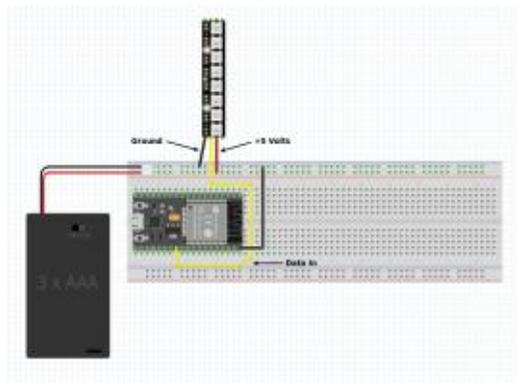

```
while True:
    for i in range(num_leds):
        for j in range(num_leds):
            pixels.set_pixel(j,abs(i+j)%10,abs(i-(j+3))%10,abs(i-(j+6))%10)
        pixels.show()
        time.sleep(0.05)
```

the data pin and pin # 38 as the ground pin. You can use whatever pins you wish.

THE CODE

RASPBERRY PI PICO

First, download the repository from the github site above. Save it to a convenient folder and unzip it. Then copy the file ws2812b.py to the pico as well as the example files.

We'll look at the flash.py example first.

```
import time
```

```
from ws2812b import ws2812b
```

Now we'll take a look at the first two lines of code. The first sets the number of LEDs that you have on your device. Be sure that it matches the number of LEDs that you really have. Many stick type displays have only 8 LEDs. Some strips have 30 to

60 on them. The ring that I have has 24 LEDs.

The second line instantiates the ws2812b class into the variable pixels. Notice that the first parameter is the number of LEDs, the second is the state machine that will be used. This is normally either a 0 or a 1. The next parameter is the GPIO pin number. The last is the delay that you want to use before resetting the LED set. It should be safe to use 0, unless you need to do a lot of processing.

```
num_leds = 24
```

```
pixels = ws2812b(num_leds, 0,
16, delay=0)
```

Next comes the brightness value. This can be useful if you are like me, and the LEDs blind you

because they are so bright. The next two lines fill every LED on your device to 10,10,10 (R,G,B values) and then the information is sent out to the LED device with the pixels.show() command.

```
pixels.brightness(30)
```

```
pixels.fill(10,10,10)
```

```
pixels.show()
```

At this point, we set up a simple loop (above) and for each LED that you have set up at the beginning of the code, which sets the color. Then it sends the data out and pauses for a little bit.

Now, the code might look fairly easy until you get to the set_pixel line. Here (below) is the actual function from the library. It makes it much easier to understand what's

```
def set_pixel(self, pixel_num, red, green, blue):
    # Adjust color values with brightnesslevel
    blue = round(blue * (self.brightness() / 255))
    red = round(red * (self.brightness() / 255))
    green = round(green * (self.brightness() / 255))

    self.pixels[pixel_num] = blue | red << 8 | green << 16
```

MICRO THIS MICRO THAT

happening.

You can see that the set_pixel function takes 4 parameters. The pixel number, and the red, green and blue values.

So going back to the loop above, we can step through the code and see the values that are being sent to the LED device. I'll show only 7 values, which should give you a good idea of what happens.

```
0  0  3  6
1  1  4  7
2  2  5  8
3  3  6  9
4  4  7  0
5  5  8  1
6  6  9  2
…
```

When we run the program, all the LEDs will blink different colors in a quickly changing pattern.

The other example program I want to point out is the fireflies.py program. It is supposed to simulate the flitting of fireflies on a summer's evening. I have to admit that it is more impressive on a 1 metre 30+ LED strip than on the 8-LED strip.

```
import time
```

```
import ws2812b
```

```
import random
```

```
numpix = 24   # Number of
NeoPixels
```

```
# Pin 16 is where NeoPixels
are connected
```

```
strip =
ws2812b.ws2812b(numpix, 0,
16)
```

Remember to set the number of LEDs that your device has as well as the GPIO pin that the device is connected to.

```
colors = [
  [232, 100, 255],   # Purple
  [200, 200, 20],    # Yellow
  [30, 200, 200],    # Blue
  [150,50,10],
  [50,200,10],
]
```

Here (top right) max_len and min_len is the maximum and minimum length of time the LEDs are on. The num_flashes refers to the number of "fireflies" that are active at one time.

This is the display function that handles the glowing and flitting of the virtual flies (which are actually beetles in real life). Shown bottom right.

```
max_len=20
min_len = 5
#pixelnum, posn in flash, flash_len, direction
flashing = []

num_flashes = 10

for i in range(num_flashes):
    pix = random.randint(0, numpix - 1)
    col = random.randint(1, len(colors) - 1)
    flash_len = random.randint(min_len, max_len)
    flashing.append([pix, colors[col], flash_len, 0, 1])

strip.fill(0,0,0)
```

```
while True:
    strip.show()
    for i in range(num_flashes):

        pix = flashing[i][0]
        brightness = (flashing[i][3]/flashing[i][2])
        colr = (int(flashing[i][1][0]*brightness),
                int(flashing[i][1][1]*brightness),
                int(flashing[i][1][2]*brightness))
        strip.set_pixel(pix, colr[0], colr[1], colr[2])

        if flashing[i][2] == flashing[i][3]:
            flashing[i][4] = -1
        if flashing[i][3] == 0 and flashing[i][4] == -1:
            pix = random.randint(0, numpix - 1)
            col = random.randint(0, len(colors) - 1)
            flash_len = random.randint(min_len, max_len)
            flashing[i] = [pix, colors[col], flash_len, 0, 1]
    flashing[i][3] = flashing[i][3] + flashing[i][4]
    time.sleep(0.005)
```

Most of the demo programs that I have found have similar loops as the above two programs. And they have a similar while loop that, when stopped, the LEDs are still active. So I wrote the following short program to clear the LEDs when I'm done with the demo. I call it led_clear.py

```
# LED Clear
```

```
from ws2812b import ws2812b
```

MICRO THIS MICRO THAT

```
num_leds = 24

pixels = ws2812b(num_leds, 0,
16, delay=0)

pixels.brightness(100)

pixels.fill(0,0,0)

pixels.show()
```

ESP32/ESP8266

My initial thought was to rewrite the flash program presented for the Pico to work on the ESP microcontrollers. However, I remembered that I had a great demo from randomnerdtutorials.com that would provide a better example of how to use the library. So, instead of flash, I present neopixel1.py (top right). The same code should run unmodified on the ESP8266, since they share the same library.

Notice that the library for the neopixel devices is named "neopixel". The variable n is the number of LEDs the device has, and the variable p is the GPIO pin that will be used.

So the variable np is the neopixel object. You need to address each pixel individually and pass the RGB values as a tuple. You can see this in the clear function presented directly below.

```
def clear():

    for i in range(n):

        np[i] = (0, 0, 0)

    np.write()
```

The next function (middle right) is called bounce, which, as you might expect, causes a single LED to move through each available position and, when it gets to the starting point, it reverses direction.

The set_color function, again, does as the name suggests. It sets all the pixels (defined in n) to the RGB color at the brightness you desire.

```
def set_color(r, g, b,
brightness):

    for i in range(n):

        np[i] =
(int(r*brightness),
int(g*brightness),
int(b*brightness))

    np.write()
```

The cycle function (bottom right) will cause a single LED to light and walk through the strip at the defined color and brightness, at

```
# ESP32 Neopixel Demo
# https://randomnerdtutorials.com/micropython-ws2812b-
addressable-rgb-leds-neopixel-esp32-esp8266/
# Modified by G.D. Walters

import machine, neopixel
import time
n = 24
p = 16

np = neopixel.NeoPixel(machine.Pin(p), n)
```

```
def bounce(r, g, b, brightness, wait):
    for i in range(4 * n):
        for j in range(n):
            np[j] = (int(r*brightness), int(g*brightness),
int(b*brightness))

        if (i // n) % 2 == 0:
            np[i % n] = (0, 0, 0)
        else:
            # pass
            which = n-1-(i%n)
            np[which] = (0, 0, 0)
        np.write()
        time.sleep_ms(wait)
```

```
def cycle(r, g, b, brightness, wait):
    for i in range(4 * n):
        for j in range(n):
            np[j] = (0, 0, 0)
        np[i % n] = (int(r*brightness), int(g*brightness),
int(b*brightness))

        np.write()
        time.sleep_ms(wait)
```

a defined speed.

The next function, wheel (next page, top right), is a support function used by the rainbow_cycle function below.

Now that we have all of the functions defined, we'll walk through each of them, one at a time. We'll print to the REPL terminal which function is being

MICRO THIS MICRO THAT

run.

```
print('Clear')

clear()

print('Set Color')

set_color(0,102,.5,230)

time.sleep(2)

clear()

print('Cycle')

cycle(0,102,230,.5,200)

print('Bounce')

bounce(0,255,250,.5,200)

clear()

time.sleep(2)

print('Rainbow Cycle')

rainbow_cycle(.5,5)

clear()

print('Program Ends')
```

It should be fairly easy to understand how the program actually works. I've included on the github repository a port of the fireflies program for the ESP called fireflies1.py. As an added benefit, I've added a small call to the Touchpad function we discussed previously that will break the

forever loop when the touchpad is detected.

All of the code and image files can be found at https://github.com/gregwa1953/FCM171_MicroThisMicroThat .

I sincerely hope that I have inspired you to get some NeoPixel type devices and play with them. After all (as if I needed to give you a reason), Christmas, Hanukkah, and the rest of the festive holiday season celebrations are less than 6 months from now. Just picture what these little devices can do for your decorations this year and how jealous your neighbors will be.

Until next time, as always; stay safe, healthy, positive and creative!

```
def wheel(pos):
  # Input a value 0 to 255 to get a color value.
  # The colors are a transition r - g - b - back to r.
  if pos < 0 or pos > 255:
    return (0, 0, 0)
  if pos < 85:
    return (255 - pos * 3, pos * 3, 0)
  if pos < 170:
    pos -= 85
    return (0, 255 - pos * 3, pos * 3)
  pos -= 170
  return (pos * 3, 0, 255 - pos * 3)

def rainbow_cycle(brightness, wait):
  for j in range(int(brightness*255)):
    for i in range(n):
      rc_index = (i * 256 // n) + j
      np[i] = wheel(rc_index & 255)
    np.write()
    time.sleep_ms(wait)
```

Greg Walters is a retired programmer living in Central Texas, USA. He has been a programmer since 1972 and in his spare time, he is an author, amateur photographer, luthier, fair musician and a pretty darn good cook. He still is the owner of RainyDaySolutions a consulting company and he spends most of his time writing articles for FCM and tutorials. His website is www.thedesignatedgeek.xyz .

Linux on Your iPad

For as low as $4.95, you can have your own personal Linux cloud computer in minutes on any device.

Written by UBports Team

Ubuntu Touch (https://ubuntu-touch.io/) is the privacy and freedom-respecting mobile operating system by UBports. Today we are happy to announce the release of Ubuntu Touch OTA-18; our eighteenth stable update to the system! OTA-18 will become available for the following supported Ubuntu Touch devices over the next week:

LG Nexus 5
OnePlus One
Fairphone 2
LG Nexus 4
BQ E5 HD Ubuntu Edition
BQ E4.5 Ubuntu Edition
Meizu MX4 Ubuntu Edition
Meizu Pro 5 Ubuntu Edition
BQ M10 (F)HD Ubuntu Edition
Nexus 7 2013 (Wi-Fi and LTE models)
Sony Xperia X
Sony Xperia X Compact
Sony Xperia X Performance
Sony Xperia XZ
Sony Xperia Z4 Tablet
Huawei Nexus 6P
OnePlus 3 and 3T
Xiaomi Redmi 4X
Google Pixel 3a
OnePlus 2

F(x)tec Pro1
Xiaomi Redmi 3s/3x/3sp (land)
Xiaomi Redmi Note 7
Xiaomi Redmi Note 7 Pro
Xiaomi Mi A2
Volla Phone
Samsung Galaxy S3 Neo+ (GT-I9301I)
Samsung Galaxy Note 4

The Pine64 PinePhone and PineTab are updated independently of the rest of these devices. The stable channel for the PinePhone and PineTab will not receive an update labeled "OTA-18".

What's new?

This release's mascot is the tangram elephant, as our primary goal with this release was keeping a steady and solid pace. We also saved some memory on all devices. Thanks for the release mascot, Joan!

This release of Ubuntu Touch is still based on Ubuntu 16.04. See the "What's new" section below for an overview of the progress of Ubuntu Touch based on Ubuntu 20.04.

Principal project packs plenty progress

Some parts of Ubuntu Touch show a strange dichotomy between being boring when described, but interesting when their impact is identified. Media-hub is one of the best examples of this. Media-hub is a service which centrally plays back audio and video content as requested by applications. Yawn.

But the features that media-hub enables are seemingly endless: a centralized service playing back media means that the platform knows exactly when content is playing and can change its behavior to suit. When a phone call comes in, your media is paused. When the call ends, your media is resumed. When you put your phone's display to sleep with a podcast playing, a centralized player can hold the system awake to prevent the audio skipping or pausing altogether during sleep. And, someday, playing back videos centrally could make today's most popular features like Picture-in-Picture playback simple.

Centralized, universal EQ settings, anyone?

It is within that exciting flurry of features and futures that Mardy realized one thing was missing... media-hub needed a facelift, and fast. The code was largely disjointed and confusing, making adding new features or fixing bugs daunting at best. This was preventing Mardy from creating an app he's always dreamed of, so he got to work. 20,526 lines of code later, the new media-hub emerged -- with better tests, a more contributor-friendly structure, and a few fixed bugs to boot.

This is one of those changes that could go completely unnoticed by someone passing over the OTA-18 release, as media-hub is completely invisible, a silent worker always in the background doing its job. But the changes that Mardy made were not easy, they improved the stability of Ubuntu Touch, and should absolutely not pass unnoticed. Thanks, Mardy!

Elite efficiency

Despite the seeming will of the internet to always have the latest and most powerful technology, we still have a lot of people using devices with just 1GB of RAM. The BQ E4.5 and E5 HD are old devices, sure, but they still work and make a lot of people happy. OTA-18 almost always feels faster than OTA-17 on the same device. Unless you know what to look for, though, it's hard to say why.

We've made Lomiri's wallpaper rendering far more efficient in this release. Without going into gory details, we've managed to coax QML's image loader into loading only one copy of your wallpaper. We've also made it scale down the wallpaper to reduce the number of pixels in RAM at any time. The amount of RAM saved by these changes varies depending on your device and whether you've set a custom wallpaper or not. If you have a custom wallpaper, you save at least 30MB of RAM on any device by upgrading to OTA-18. A larger-resolution wallpaper will use far less RAM on OTA-18 than on OTA-17, but the savings are greater for devices with a low screen resolution, so on an older device like the E4.5 with the default wallpaper, expect to save about 60MB of RAM.

You'll also see better performance across the board. With a smaller background image, it's easier for Qt to put together the scene that is displayed on your phone. Faster scene rendering means higher frame-rates no matter what is happening on screen.

You don't need to do anything special to take advantage of OTA-18's improved performance. Just download and install the update!

Annihilated annoyances

There are a lot of fixes that cost little in code but will have a huge impact in practice. We have listed out those here:
• The on-screen keyboard [OSK] now appears automatically when a new tab is opened in Morph Browser.
• Pressing Control+Alt+T now opens the terminal app.
• There is a degree (°) symbol on the on-screen keyboard.
• Stickers have been added to the built-in messaging app. If your conversations often involve sending a set of images that capture your reaction, this feature is for you!
• Alarms now snooze from the time they were snoozed rather than from the start of the alarm. They also snooze when you miss them, instead of dismissing.
• Call audio was fixed on the Google Pixel 2

What's next?

Our previous posts have alluded to a slowing down of Ubuntu Touch development on Xenial as we prepare a release of Ubuntu Touch based on Ubuntu 20.04. It seems that the fabled slowdown was if anything underestimated.

It is true that the small team of people who know the internals of Ubuntu Touch has been preoccupied with things other than OTA-18. Ratchanan has been focused on making Lomiri, the infrastructure around it, and the keyboard run under systemd on Ubuntu 20.04; on creating images of UT based on 20.04; and on too many other tasks to count. Rodney has focused on other dependencies of Lomiri and the Ubuntu Touch experience, building the fingerprint recognition service, URL Dispatcher, content-hub, and the lib-cpp projects. He's also replaced the component of the Click packaging system which allows users without `sudo` privileges to install packages. Alfred has created the Ubuntu Touch Platform Development Kit, which makes it far easier to set up an environment for developing the Ubuntu Touch platform than ever before. All of this work means that people curious about platform development and app developers can start to get in on Ubuntu Touch based on Ubuntu 20.04 right now. It is definitely true that the more work we get done on 20.04, the faster we can go.

Even with so many people central to Ubuntu Touch being so busy, the OTA-18 release is still larger and more impactful than OTA-17. Despite OTA-18's development cycle being just 4 days longer, we've managed to change thousands more lines of code.

OTA-19 will be another release based on 16.04, but the progress to 20.04 will be amazing. Let's get started.
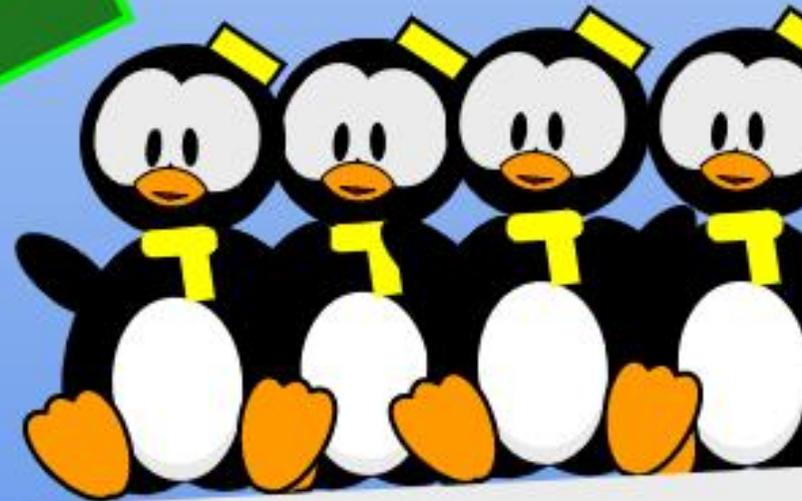
MY OPINION

Written by ...

NOTHING TO SEE HERE.

Submit your stories and opinions to:
ronnie@fullcirclemagazine.org.

GUIDELINES

The single rule for an article is that **it must somehow be linked to Ubuntu or one of the many derivatives of Ubuntu (Kubuntu, Xubuntu, Lubuntu, etc).**

RULES

• There is no word limit for articles, but be advised that long articles may be split across several issues.

• For advice, please refer to the **Official Full Circle Style Guide**: http://bit.ly/fcmwriting

• Write your article in whichever software you choose, I would recommend LibreOffice, but most importantly - **PLEASE SPELL AND GRAMMAR CHECK IT!**

• In your article, please indicate where you would like a particular image to be placed by indicating the image name in a new paragraph or by embedding the image in the ODT (Open Office) document.

• Images should be JPG, no wider than 800 pixels, and use low compression.

• Do <u>not</u> use tables or any type of **bold** or *italic* formatting.

If you are writing a review, please follow these guidelines :

When you are ready to submit your article please email it to: articles@fullcirclemagazine.org

TRANSLATIONS

If you would like to translate Full Circle into your native language please send an email to ronnie@fullcirclemagazine.org and we will either put you in touch with an existing team, or give you access to the raw text to translate from. With a completed PDF, you will be able to upload your file to the main Full Circle site.

REVIEWS

GAMES/APPLICATIONS
When reviewing games/applications please state clearly:

• title of the game
• who makes the game
• is it free, or a paid download?
• where to get it from (give download/homepage URL)
• is it Linux native, or did you use Wine?
• your marks out of five
• a summary with positive and negative points

HARDWARE
When reviewing hardware please state clearly:

• make and model of the hardware
• what category would you put this hardware into?
• any glitches that you may have had while using the hardware?
• easy to get the hardware working in Linux?
• did you have to use Windows drivers?
• marks out of five
• a summary with positive and negative points

You <u>don't</u> need to be an expert to write an article - write about the games, applications and hardware that you use every day.

The Kubuntu development team has continued its refinements with Kubuntu 21.04, introduced on 22 April, 2021. This is the middle release of this development cycle leading to Kubuntu 22.04 LTS, which is due out in April 2022. Kubuntu 21.04 marks the 33rd Kubuntu release and the 12th one with the KDE project's Plasma 5 desktop which is based on the Qt toolkit, so this is a very mature software project. As a "standard" release, Kubuntu 21.04 is supported for just nine months, until January, 2022. One more standard release, Kubuntu 21.10, is expected on 14 October 2021, prior to the long term support (LTS) release in the spring of 2022.

Booting It Up

I downloaded the Kubuntu 21.04 ISO file via bittorrent and did a SHA256 sum check on it to verify integrity. Then I used UNetbootin to write it to a USB drive for testing.

Booting to the USB stick in 21.04 no longer results in an overt file system check. Instead it is run silently in the background during the boot process, according to Rik Mills, a Kubuntu and Ubuntu developer. I presumed my USB write was good, as I did not see any notifications to the contrary.

New

This version of Kubuntu uses the Plasma 5.21 desktop, which is based on the Qt 5.15.2 toolkit. This new version of Plasma 5 introduced some small refinements and seems to run more smoothly, with better menu functionality.

The kernel used is Linux 5.11 and that provides support for a new generation of hardware.

Kubuntu 21.04 has a new default wallpaper. This one is called Milky Way by Rucker, and it is very dark in color. It looks like it is intended to compliment the dark themes that seem to be popular among desktop users these days. This release has a total of 29 wallpapers provided, including many from recent Kubuntu releases, so there are lots to choose from if this one is not to your taste.

Even though Ubuntu 21.04 is now using Wayland by default, Kubuntu and the other flavors have not followed suit yet, allowing the Ubuntu developers the opportunity to iron out the bugs first. Kubuntu 21.04 does offer an optional Wayland session that can be selected on boot-up. The Kubuntu 21.04 release notes say, "the Wayland session is provided for testing and evaluation only, and is not supported", so you have been warned.

There is only one application change in this release, other than just bumping the application version numbers to KDE Applications 20.12. The new Plasma System Monitor is added, while leaving the previous KsysGuard system monitor in place as well. I tried them both out and they seem surprisingly similar, so I suspect that KsysGuard will disappear from the default installation by the time the LTS release arrives.
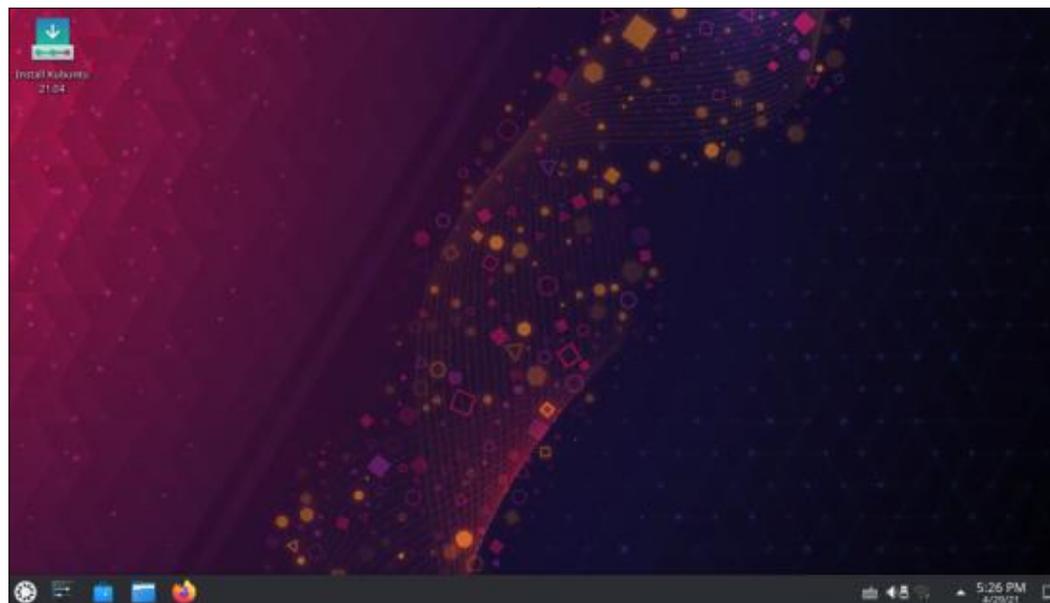
REVIEW

In general, the overall changes are small and cautious in Kubuntu 21.04, just tweaks really.

APPLICATIONS

Some of the applications included with Kubuntu 21.04 are:
• Ark 20.12.3 archive manager
• Discover 2.1.2 software store*
• Dolphin 20.12.3 file manager
• Elisa 20.12.3 music player
• Firefox 87.0 web browser
• Gwenview 20.12.3 image viewer
• Kate 20.12.3 text editor
• Kcalc 20.12.3 calculator
• KDE Partition Manager 20.12.3 partition editor
• Konsole 20.12.3 terminal emulator
• Kmahjongg 20.12.3 game
• Kmines 20.12.3 game

• Konversation 20.12.3 IRC client
• Kpatience 20.12.3 game
• Ksudoku 20.12.3 game
• KsysGuard 5.21.4 system monitor
• Ktorrent 20.12.3 bit torrent client
• LibreOffice 7.2.2 office suite, less only LibreOffice Base database
• Muon 5.8.0 package manager*
• Okular 20.12.3 PDF viewer
• Plasma System Monitor 5.21.4 system monitor
• PulseAudio 14.2 audio controller
• Skanlite 2.2.0 scanning utility*
• Spectacle 20.12.3 screenshot tool
• Startup Disk Creator 0.3.9 (usb-creator-kde) USB ISO writer
• Thunderbird 78.8.1 email client
• VLC 3.0.12 media player

* indicates same application version as used in Kubuntu 20.10.

As can be seen from the lack of asterisks, most of the applications included are updated versions from KDE Applications 20.12, with very few hold-overs from Kubuntu 20.10.

Firefox continues as the default web browser. It is a really good browser but continues in Kubuntu 21.04 as the sole GTK toolkit application in a Qt operating system. At one time, there was hope that Falkon, the Qt-based browser that was brought into the KDE project four years ago in 2017, would mature into a viable Firefox replacement but, with no releases in over two years, it is looking like abandonware.

Like previous recent releases, Kubuntu 21.04 does not include a webcam application, image editor, video editor, or CD/DVD burning application. There are choices for all of those that can easily be installed from the repositories, if needed.

The included Kate 20.12.3 text editor is worth mentioning, as it has many syntax highlighting color scheme options and now has spell-checking right out-of-the-box, which is always useful.

The LibreOffice 7.1.2 office suite is included as well and is complete, less only the LibreOffice Base database application. Base is one of the least-used LibreOffice components but can be installed if desired.
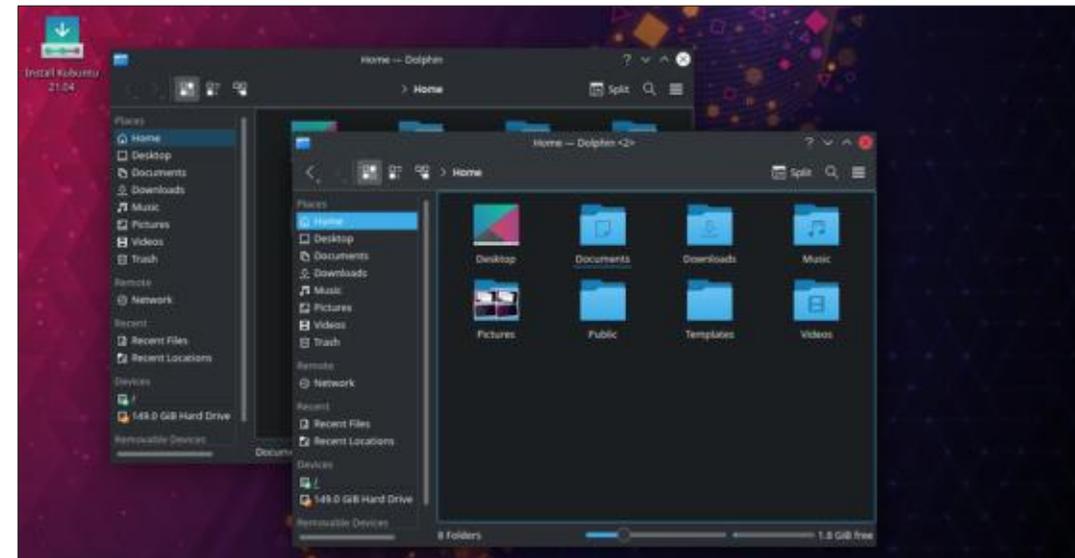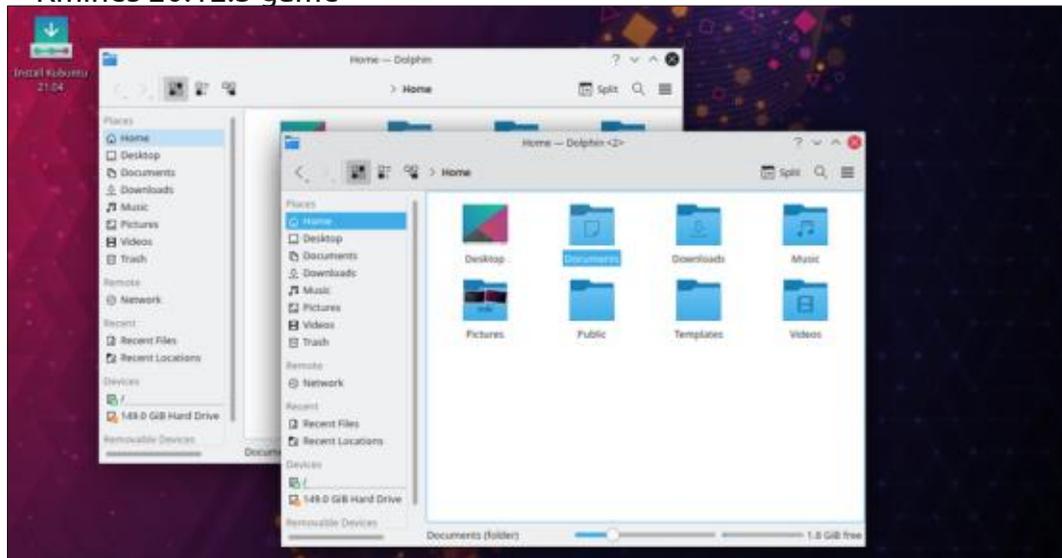
SETTINGS

Kubuntu has always been known for its high degree of customization available. In fact, that is probably one of the main reasons that it has such a dedicated following among its user base – it can be made to look and work almost any way that you like.

In a lot of ways, the Kubuntu philosophy is the diametric opposite of that of Ubuntu which offers only a very minimal selection of user configuration options.

Kubuntu has so many settings and options that running through them and setting your desktop up can be a bit daunting. For example, in Kubuntu 21.04, there are:
• Four global themes: Breeze, Breeze Dark, Breeze Twilight, and Kubuntu
• Six Plasma styles: Air, Breeze, Breeze Dark, Breeze Light, Kubuntu, and Oxygen
• Two window decoration themes: Breeze and Plastik
• Six colour schemes: Breeze, Breeze Dark, Breeze High Contrast, Breeze Light, Oxygen, and Oxygen Cold
• Eight icon sets: Adwaita, Breeze,

Breeze Dark, Humanity, Humanity-Dark, LoginIcons, Ubuntu-Mono-Dark, and Ubuntu-MonoLight
• Eight cursor styles: Breeze, Breeze Light, KDE Classic,Oxygen Black, Oxygen Blue, Oxygen White, Oxygen Yellow, and Oxygen Zion

And these are just the installed options. Many of the settings pages offer one-button downloading of large numbers of more choices, too.

Kubuntu 21.04 has 68 widgets installed that can be placed on the desktop, which is two more than in the last release. Widgets are small programs that add desktop functionality. Some examples that I tested included an analog clock, a simple menu, and an application

launcher. The latter is a full-screen main menu, much like the one found on Ubuntu. Since the Kubuntu 21.04 main menu now works better than in the past, these alternative menus may not prove as necessary as they once were.

Overall, I can see why Kubuntu users would want to avoid doing a "fresh installation" and prefer to just upgrade. With all the settings available, it would be a challenge to duplicate your past installation, even with good notes, plus the options change from release to release, also.

CONCLUSIONS

Kubuntu 21.04 is another solid release. Like most Ubuntu-based distributions, the majority of Kubuntu users stick to the LTS releases and avoid the standard releases in between, due to the short, nine-month support periods, unless there is a really compelling reason to upgrade. Kubuntu 21.04's desktop upgrades, especially the improved menus, might just convince some users that it is worth the switch.

With two releases now out in this development cycle, it seems clear that we can expect only small and incremental changes on the way to Kubuntu 22.04 LTS, which is a good thing for the users, really. Most Kubuntu users don't seem to want large scale changes. There is one more standard release to go, but I am expecting only refinements to ensure that the LTS that follows is as smooth as possible.
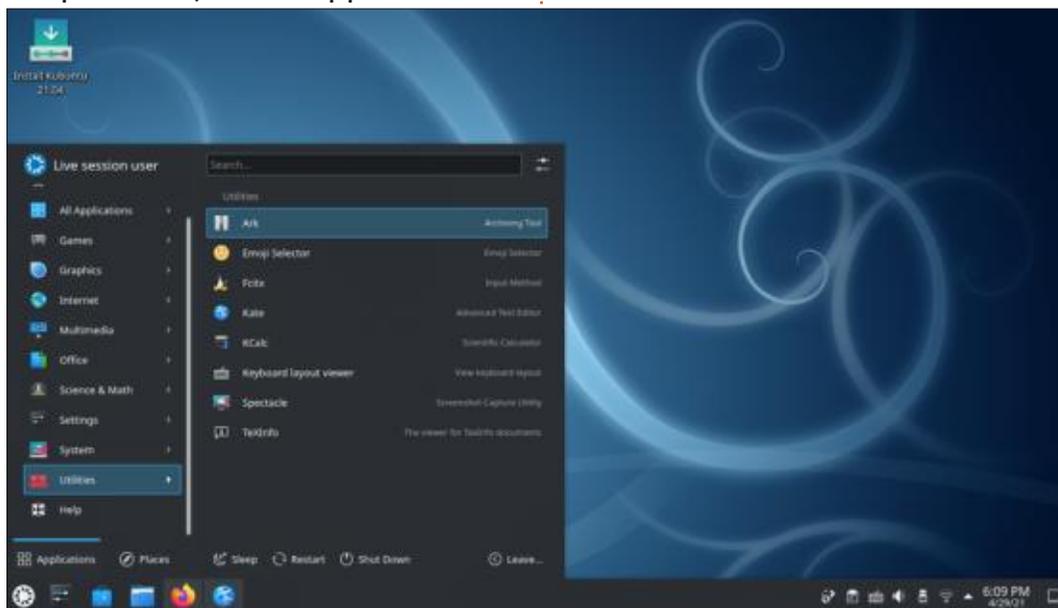
Adam Hunt started using Ubuntu in 2007 and has used Lubuntu since 2010. He lives in Ottawa, Ontario, Canada, in a house with no Windows.

Website: https://codelite.org/

Price: Free!

Blurb: "*CodeLite is an open source, free, cross platform IDE, specialized in C, C++, PHP and JavaScript (mainly for backend developers using Node.js) programming languages which runs best on all major Platforms (OSX, Windows and Linux)*"

At the time of writing, we are on version 14.0.2.

Why had I never heard of CodeLite before?

I had to see why this awesome IDE was not a lot more popular... Maybe because it says it is mainly for back-end developers using Node.js? Searching the term "codellite issue" brought up quite a few issues. This is NOT a BAD THING (they all have issues). It shows that there are people using it.

I fired it up and a helpful wizard helped me configure the basics. I

created a new Workspace and ran through the five questions that you need to get your project off the ground. Just like Code::Blocks, it created a nice 'hello world' console app that I could use as a base. Just to test, I changed the default gnome-terminal to the built-in terminal. It required a restart, but nothing was lost. Building and running the project was a breeze. The "built-in terminal" is actually another floating terminal, so nothing lost or gained there. The only plus I could see was that the theming integrated well with the

CodeLite terminal, but if you have configured your OS terminal the way you like it, I see no reason to switch. I would have liked it to be part of the bottom bar, but I can also see how this is useful.

Let me get this out of the way from the start, I am in no way an IDE power user or coding guru, but even for someone with my meagre skills, this IDE was not only handy, but easy to use and easy on my eyes at night (it is also very light on resources!).

CodeLite is arranged in a standard, familiar, maybe boring, three-pane layout, with a menu bar along the top. You also have the option to close any of the panes you do not use and have one large programming pane only! For distraction-free programming, ctrl+F11 will maximize the screen real estate and hide the menus. All the menu items are standard – or seem standard – until you reach the 'plugin' menu. This not only lets you manage your plugins, but also each plugin's settings. The settings menu proudly displays the colors and fonts menu item in the drop-down, so there is no faffing in the tools and settings when you need to switch from night to day. In that same menu, you will find the keyboard shortcuts, all the handy things in one place.

The code window is unremarkable, the way I like it - and small things, like folding functions, are easy to spot. Places where the mouse would be in the way, the mouse direction swops! They say the devil is in the details, and this IDE sure pays attention to the
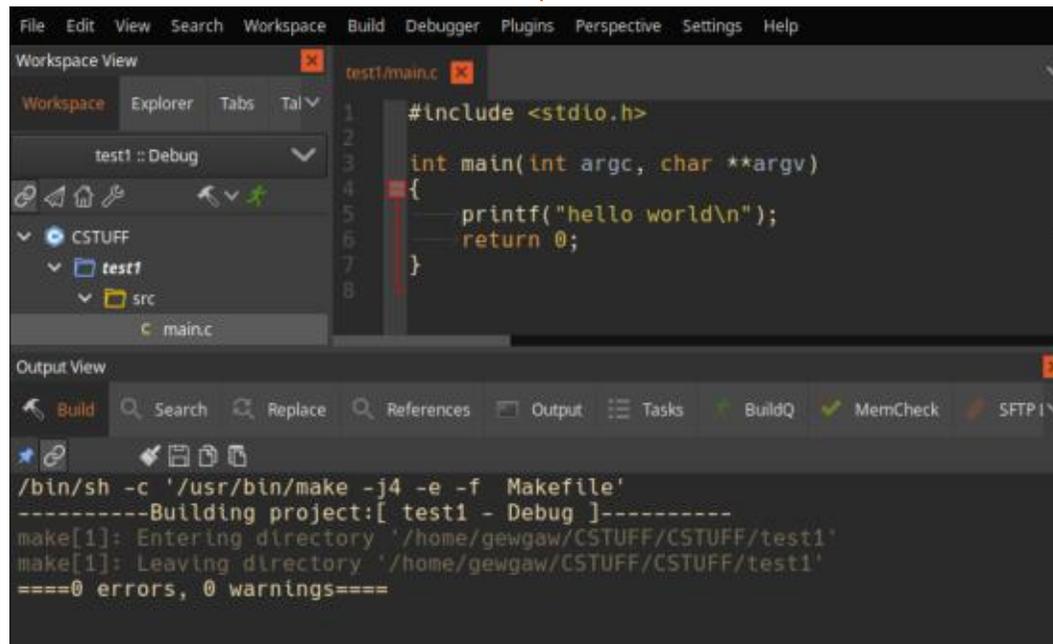
details. (I would like to show it to you, but for some reason, if I add "include pointer" in the screenshot tool, it flips the pointer back.) Right-clicking anywhere in the application, will bring up that specific part's right-click menu. There is no 'global' right-click menu, so just be aware of that. Clicking in different bands and sections is recommended, so you can get familiar with what-is-what. Notifications appear under your code, and do not interfere with what you are doing.

The 'Output View' contains a myriad of tabs I do not use, but I am sure they are important as you see those same tabs in other editors, like Code::Blocks. At my level, "build" and "output" are the only two I use (the same with the tabs in the workspace view).

Some things that may trip up novices (I include these here as the negative reviews about this fantastic little IDE usually include these):
• When you have multiple projects, using the compile-and-run will compile and run the last project you worked on. You need to double-click the new project name until it highlights, before it is the 'active' project.
• Installing on Ubuntu? Remember to install build-essential first.
• Workspaces are important and can be shared. That's right, more than one person can work in a workspace!
• Just opening random C-code will not compile – see our first tip.
• Plugins that do not work – use the help -> check for updates.
• If you use it on screens with shoddy resolutions, ctrl +0 will reset the zoom level.
• An easy way to back up your project is to simply duplicate the tab and save it somewhere else.

On to what I did not like: I could not see an easy way to clear my recent files list. This is important when you work on files on other mediums, as the projects on my flash drive now eclipse my projects on my local drive. For some reason, the ctrl+shift+up/down did not work for me to move a line up or down (but - that may be due to the way I have my OS shortcut keys defined.)

Overall, I cannot complain about anything in this free and open source IDE. I just wish I knew about it before I ran headlong into eclipse, which is so busy, it makes me tired just looking at it. I suspect the biggest issue with the application is that it does not have a 'shoutbox', hype or community blow up. Countless websites and FOSS pages sing praise to Atom code editor that is an electron application, but none praise this brave little IDE that "just works" and does not carry oodles of bloat around. According to HTOP, Codelite used less CPU and memory than my music player while I was using it. Only Sublime text and notepadqq were lighter on resources. Given, I write only small CLI programs on it, but it compiles pretty fast and I never feel like I am waiting for it.

This is a solid FOSS project I can recommend with a smile. What are you waiting for, give it a try!

As usual, if you think I was unfair or made a mistake, feel free to email us on: misc@fullcirclemagazine.org
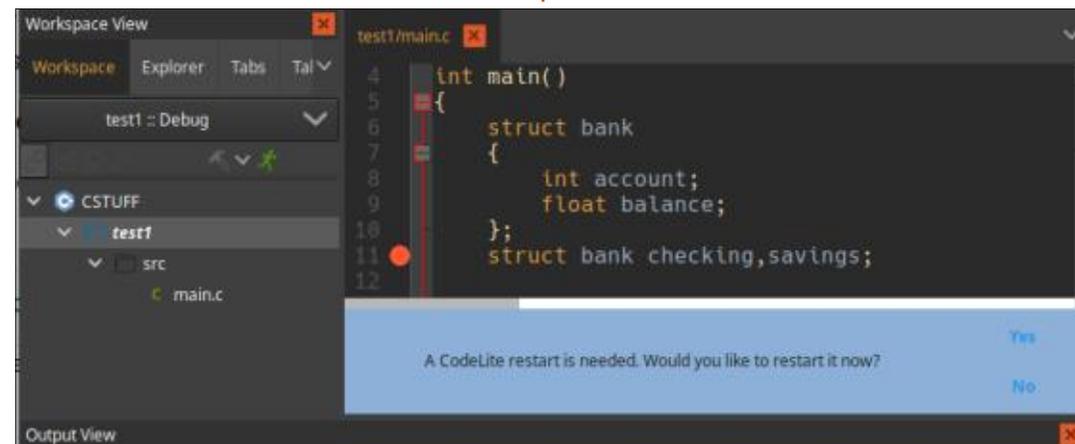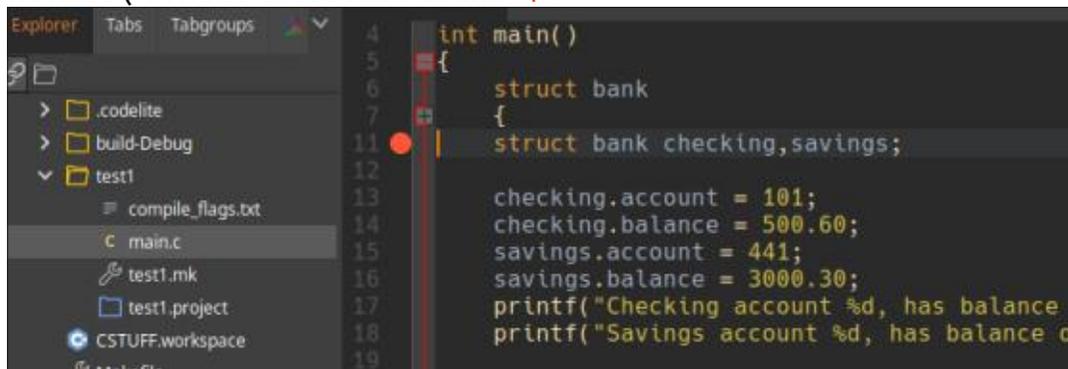
A CodeLite restart is needed. Would you like to restart it now?

COOKING WITH UBUNTU

First thank you for your magazine, it is always a pleasant moment to read it.

Regarding Richard Adams, and his articles about "Cooking with Ubuntu", thanks to him for the "tour" of recipe apps he is providing, I also did some research to find the best software for my recipes.

I agree Gnome Recipes is promising, but obviously the lack of import functions makes this choice impossible for me, as I would have hundreds of personal recipes to re-enter. Also, development looks very slow, and I could not find any milestone about adding an import function. I think this is not in the scope of the app.

Today, I use Gourmet 0.17.4, which is still usable with Ubuntu 20.04, even if you do need to manually install python 2.7 and download some packages from previous Ubuntu releases, see (https://askubuntu.com/questions/1230167/can-gourmet-still-run-on-ubuntu-20-04).

IMO, Gourmet is the best solution, because it has all functions, especially import and export with multiple formats. This is important, as I can export my recipes to .mcb (my cookbook) format, and import the file with the Cookmate app on Android (for example). That way, I can cook while travelling outside of my house, at a friend's home typically, which is mandatory for me.

Lasty, there is a Gourmet migration project to python 3. There is already a flatpak package available, and it seems usable, see (https://github.com/kirienko/gourmet/blob/master/INSTALL.md). I just tried it in a VM, and it looks good so far.

Pascal

FULL CIRCLE NEEDS YOU!

Without reader input **Full Circle** would be an empty PDF file (which I don't think many people would find particularly interesting). We are always looking for articles, reviews, anything! Even small things like letters and desktop screens help fill the magazine.

See the article **Writing for Full Circle** in this issue to read our basic guidelines.

Have a look at the last page of any issue to get the details of where to send your contributions.

Q&A

Compiled by EriktheUnready

If you have a Linux question, email it to: questions@fullcirclemagazine.org, and Erik will answer them in a future issue. Please include as much information as you can about your query.

Welcome back to another edition of Questions and Answers! In this section, we will endeavour to answer your Ubuntu questions. Be sure to add details of the version of your operating system and your hardware. I will try to remove any personally identifiable strings from questions, but it is best not to include things like serial numbers, UUIDs or IP addresses. If your question does not appear immediately, it is just because there is such a lot, and I do them, first-come-first-served.

I called Dell support on behalf of one of my clients. Dell support is 50/50 – sometimes you get good service, sometimes not. The client rents thousands of machines worldwide, and works very closely with Dell to have a unified IT infrastructure. As usual, I tried to speed up the process by telling the poor first-liner what I had done. However, this was one of those idiots who was going to treat me like a home user. I had to go through everything manually with him on the phone. (What testing VGA modes has to do with a faulty USB-C port is beyond me). I tried to explain to him that we have tested by process of elimination and come to the conclusion that the laptop's USB C port was indeed faulty. However, this guy was not budging. He wanted us to change drivers, etc, which we (the client) are not allowed to as it is controlled via the client's global IT in combination with Dell. There is monitoring software that will remove the PC from the network if we change anything – as it will be non-compliant. Here we got stuck, they refused to help until I updated the drivers. The thing is, we have 100's of similar laptops with everything identical, so chances of the driver causing one machine to be faulty are tiny. Then the guy asked for a photo of the port as he was sure it was damaged. This was not my first rodeo, so I sent him a picture of mine – my laptop being less than a month old and the port never used. (The other laptop was about 2 years old). Dell does not want to waste resources on older laptops, and while I understand that, I have a 3-year on-site warranty that is paid for. The reply was that the port was physically damaged and they would void the warranty. This is where I lost it and became a Karen. His manager tried to tell me that it was faulty too at first, (trying to cover) until I told him there is NOTHING wrong with it as it was mine and unused. Only when they realised they had been duped and that I was on to them from the start, did they want to see the other one, which, by-the-way, looked the same. One can't really see from a photo if a USB C port is damaged internally, only externally, like bent or broken edges or centrepiece. I wonder how many clients they have done this to, to be so confident in trying to BS me into accepting some first liner's opinion. They tried hither and thither, and tried to pull EULA, etc, on me, but I did not give in. Don't let someone try to tell you because you run Ubuntu your warranty is void or they can't assist. BS back, tell them you installed Windows and your driver is now the version they want and it is still faulty. If you are 100% confident in your diagnosis, stick to your guns. (BTW, they replaced the board and it has been OK ever since, but do check that too, Dell does not always send new parts, a lot of them are "refurbished" and may have other issues).

_____

Q : I installed Balena Etcher from AppImage onto my Ubuntu 18.04. It worked fine. I needed some space on my SSD, so I deleted the AppImage itself. Now it won't work. Do you want to tell me I need to keep the AppImage always?

A : Pretty much; you see, AppImages contain all the files the application needs, should your system not have them. This removes the need for installing dependencies, but at the cost of space.

Q : I installed Ubuntu 21.04. Everything works great, except that I can't take screenshots. It looks like crap. I have done an apt update, so everything is latest. Should I try the PPA version? I don't really want PPA's.

A : None of that is relevant, You need to switch to X and not Wayland, to use screen grabbing software.

Q : Okay, here goes, I need a good working auto clicker for linux ubuntu, and I can't seem to find a working one. I need it for the cookie clicker to click while I am at school. I recently switched from Windows as my computer can't handle Windows 10.

A : Have a look at xdotool. We had a few issues featuring a tutorial a while back. https://www.semicomplete.com/projects/xdotool/

Q : What means this --> Buffer I/O error on dev sda, logical block 0, async page read
 sd 0:2:0:0: [sda] tag#969 FAILED Result:
hostbyte=DID_NO_CONNECT
driverbyte=DRIVER_OK
cmd_age=0s
 sd 0:2:0:0: [sda] tag#969 CDB:
Read(10) 28 00 00 00 00 00 00 00 08 00

A : I/O error on sd usually means your CD or DVD is bad.

However, it could be a bad CD/DVD ROM or a bad cable. You would need to eliminate those first, then ask me again.

Q : Hi guy. My dmesg output is filled with this weird error. I know UFW is my firewall and I was wondering if I am being hacked? The output is - [UFW BLOCK] IN=wlp3s0 OUT=
MAC=01:00:5e:00:00:01:50:e0:36:08
:85:e0:08:00 SRC=192.168.1.1
DST=224.0.0.1 LEN=36 TOS=0x00
PREC=0xC0 TTL=1 ID=40140
PROTO=2
over and over.

A : 224.0.0.1 is a broadcast address. It could be some chop scanning, but I am going to assume it is some equipment on your ISP side. Swop your ISP router for your own and see if the problem persists?

Q : My new laptop has a 256GB SSD and a 1TB HDD. How can I see the combined free space on these drives instead of running a df -h on both and using a mate calculator to tally the result? Am I making sense, do you understand what I am trying to do?

A : Yup. If you run a man df, you will see a --total (two dashes, no space). The description is: "--total elide all entries insignificant to available space, and produce a grand total" So, df --total -h /dev/sd* should get you what you need.

Q : I have recently discovered Synthwave thanks to Rocket league, and it is so retro and yet awesome. YouTube always opens in a large size in Firefox. It is costing me a fortune in mobile data, so I want it to always open small. How can I do that? Ubuntu 18.04, HP Core2Duo 2.0, 500GB and 4GB Ram.

A : You need an app for that. (See: https://linuxhint.com/best_youtube_apps_linux/) Check your music player, it may have an integrated option for playing directly from Youtube. You can also use youtube-dl to grab only the audio and play it back at your leisure.

Q : I am having issues with the builtin file roller, when I try to unzip certain files. Sometimes it says that the file is not a compressed file and sometimes it just spits out an error at me. When the files are large or I want to see the contents of an ISO file, my Xubuntu just seems to freeze. Do you know why it does this?

A : Nope, no idea, but I would suggest installing 7zip, winrar, etc, and maybe change the front end to something like engrampa instead of file roller. Let me know how it goes.

Q : I am sorta stumped here. One of my assignments says that I need to make a note of the version with lsb_release. However, when I run it on Ubuntu 18, I get an error: "shawn@powertop ~ lsb_release -a No LSB modules are available." How can I fix it?

A : Truth be told, I have not used that on Ubuntu based stuff, only Red Hat based, so before I help you wrongly, you can run inxi -S or use neofetch or even just cat /etc/os-release or hostnamectl - there are many ways to skin that cat in Linux.

Q : My Wifi router is in the hallway of the house, on the opposite side of where my room is. I suspect that the distance is the maximum

for the zyxel. I have 2 laptops, one with Windows and one with Ubuntu. The Windows one will disconnect, but the Ubuntu one does not have a WM installed; it's emacs all the way down. How can I check the wifi strength that Ubuntu gets?

A : It depends on how modern your Ubuntu is, but try nmcli. If unsure, first try this: nmcli networking connectivity - that should tell you if your WiFi connection is poor or disconnected.

Q : I managed to get a first gen i5 machine from my friend in <Removed>. I was using a core 2 Duo E7550 before on MSI motherboard. There, all was okay in Ubuntu. This new one does not list the soundcard. It is onboard Dell. Why does Ubuntu not see the Dell sound?

A : My guess would be that you have it turned off in the BIOS. It has been a really long time since I have seen Ubuntu not detecting a sound card on a Dell motherboard. The other option is that it may be faulty. You can test with another card, or even a USB sound card, and see if it is detected and if there are

any resource clashes with the on-board one.

Q : So I read about hardware level encryption for the new SSDs. I am using LUKS atm, and I was wondering if I could tell if it is activated in Ubuntu. My machine is an i7 Dell and my SSD is Adata.

A : That is not such a straightforward question. Did you enable TPM in your BIOS; does your computer ask you for a password when you turn it on, before it boots? Did you use any manufacturer software to do it? If not, you are just using software encryption.

Q : I was following the instructions on here: http://www.compholio.com/netflix-desktop/ but there is an error this ^ E: The repository 'http://ppa.launchpad.net/ehoover/compholio/ubuntu focal Release' does not have a Release file. N: Updating from such a repository can't be done securely, and is therefore disabled by default. N: See apt-secure(8) manpage for repository creation and user configuration details.

A : If you look at the launchpad page here: https://launchpad.net/~ehoover/+archive/ubuntu/compholio - you will see it is archived and goes up to only Ubuntu 15.10. So you should get a 404 – page not found, as it does not exist for Ubuntu 20.04.

Q : Can I take my Ubuntu installed Raspberry PI SD card and insert it in a newer Raspberry PI, without breaking my installation? I have done a lot on that little 16GB card, and I don't want to do it all over from the start.

A : Linux is a lot more forgiving than Windows and will probably work with 80% certainty. However, it is a risk, therefore I would rather advise you to clone the SD card and try it. That way, you have nothing to lose.

Q : I have just finished customizing Xubuntu 20.04 the way I want it. However, I have noticed that when I copy and paste to the mousepad, the "-" minus is missing. Any idea as to why?

A : Change the default font to something else. Not all font packs are 100% complete.

Q : This is going to sound stupid, but I clicked the view menu away by mistake in files on Xubuntu. I have no idea how to get it back. I held down alt and dragged my finger over most of the keys and nothing. I am a total beginner here and I don't have the foggiest idea what else to do, right-clicking in most parts does nothing, and the icon in the top-left does not give me the option. Googling for, like forever, got me zilch. What else can I do?

A : You said Xubuntu, but you said "files" - this is a bit confusing. Xubuntu uses Thunar, vanilla Ubuntu uses Nautilus, aka Files. You can just press CTRL+M to get the menu showing permanently or just pressing F10 will pop it into existence temporarily, should it be Thunar in Xubuntu.

Q : In Windows, I can see the history of Wifi networks I have connected to in the past. In my laptop I can click on the Wifi icon and go to connections and see

Q&A

them there. This I am Okay with. I have built a headless Wifi scanner with a Raspberry Pi zero from a tutorial, and I was wondering if I could see what open networks I have connected to, but from the command line. It has no DE. I can install one, but I don't really want to.

Oh yes, it is Ubuntu 18.04 minimal on there, I had Ubuntu Mate on, but it was slow.

A: You can try:

```
nmcli connection show
```

it should show you what you need. It also has a small help:

```
nmcli -h
```

Q: Smorg, [10.03.21 11:22] I know it's not 100% Ubuntu, but I want to put it in my VM on Ubuntu, so it may count. Why would anyone want to use Openindiana? They say it is rock-solid.

A: Legacy, mostly, I'd say. People like to have at home what they use at work. Thing is, commercial software for it is like hens' teeth, and that is what is important at the

end of the day; if nothing new runs on it any more, would you want it? Linux has been "Rock-solid" for me moreso than Windows 10, but that does not say much. Also remember, it is based on Solaris, proprietary OS, meaning open source stuff for it is also minimal. But go ahead and install it, you may like it.

Q: How can I create a file containing a minus? Like touch -123? I was trying to make a file my little brother can't just delete.

A: Don't, in Linux a 'minus' is used to add options to commands. You may end up with something you did not intend. Use user/file permissions instead. Instead of trying to be clever, be smart.

Q: I am new to this home automation thing. Will my drive always be sda2? Why not 1?

A: No, Ubuntu is not Windows, to see your disks and partitions use:

```
sudo fdisk -l
```

to list them all. The number is the

number of the partition, not the drive.

Q: Just a quick question about LibreOffice in Ubuntu 18 vs 20. When I paste via "paste special" and "unformatted text", something odd happens. The character fonts will assume my document fonts, but the spacing or line width will not. I noticed this while editing a document on my wife's Ubuntu 20. It is very annoying. I can delete the line and press enter for a new line and the spacing – by this I mean the space an empty line takes up, will be completely different. This throws off everything else you want to line up. Then two lines further, everything is fine again. Like I don't have enough grey hairs already.

A: What you can try is select everything CTRL+A, click on 'Default style' top left, then set your fonts and font size again. When copying and pasting from browsers, people try to insert hidden nonsense, designed to make your copy/paste life difficult. When using copy/paste from a browser, you can turn off javascript, that may help. You can also make sure you do not select any empty spaces. You

can also copy/paste to a plain text editor, and there select the text you want and copy/paste that.

Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

Website: https://huntdown.com/

Price: $19.99 USD

Blurb: "*In the mayhem-filled streets of the future, where criminal gangs rule and cops fear to tread, only the bounty hunters can free the city from the corrupt fist of felony. Lay waste to the criminal underworld and make a killing in this hard-boiled action comedy arcade shooter.*"

Before we do anything, I'd like to point out that the game script is 345MB and the soundtrack is 800MB.

While I could rag on about the lame title, I won't. The game loop is simple, a run-and-gun a lá 90's arcade. Skin it slightly differently and you could be playing Robocop from 1989. Heh, one of the characters IS a robot. But I will play the dude with the builtin harmonica – I want to imagine mirrorshades built in. This throwback to old school gaming even has punks as the enemy. Yes,

punks. One thing I can say without hesitation is that the game is easy to grasp and play; no manual required. This could easily be an Atari ST or Amiga game.

There are a lot of cutscenes and story woven in, bully for Easy trigger games. With Cyberpunk 2077 out in the wild, this Cyberpunk dystopia stays with the flavour of the year. That doesn't mean there is a lot of story here, but it has a lot more than most I have played. At first, the graphics are amazing, bang on the money for an old-timey arcade game. But if

you play it for longer than five minutes, you realise that it is very generic. By that statement, I mean that they don't really stand out, not that they are bad. In fact, the animation is brilliant. The backgrounds are great, the sprites are great, just your main character is a bit unremarkable. Games like this would be - Contra, Army Moves, or Midnight Resistance. Earthworm Jim or Ruff n Tumble, or Metal Slug, on the other hand, had really interesting lead characters. Like a story in a book, we need to identify or get attached to our main character to maximize our

enjoyment.

As a bounty hunter, you battle through waves of well animated cannon fodder... to get to the sub-boss... to get to the boss, then to repeat the recipe. There are some interesting weapon pick-ups as well as hidden loot. Though this is a game, it feels more like a Hollywood action flick. (Queue Robocop again.) Gunfire, explosions and inept cops help immerse you into the movie... erm game.

Movement is slick and the controls are responsive and on the money. When you duck into cover, there is no doubt, and when you shoot a punk it is immediate. When the screen is full of enemies, not once did it feel like the game was going to slow down, even on my potato laptop. I did say there was a story, and the dialogue is very cheesy. This is because there is some humour built in, like the ambulance-thingy checkpoint (Tony's surgery and repairs) that will kick you out on your backside after you die. Some things are just stupid though, like the name Anna
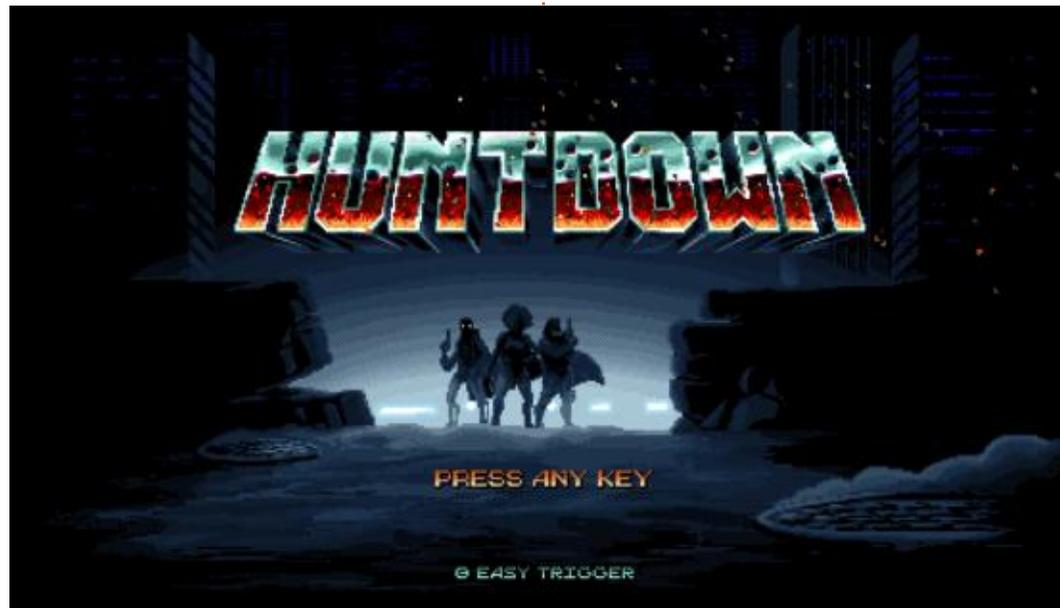
Conda. Speaking of which, some of the references have been changed slightly and won't make sense to new players; for instance, Duke Nukem's "Hail to the King baby!" is "Hail to the Queen, baby!"

The music is great. It does not get in the way of the game, but it adds to the aesthetic. Hey, I have the soundtrack playing while I am writing this. I will probably turn the volume down slightly and go to sleep with it playing.

But wait, there's more... (as they say in the commercials). The SFX and voice acting are also A+ quality. Yes, this alone makes the game wear a red peak cap. Make the

game sound great again. So don't be shy; pump up the volume. There has obviously been a lot of effort put into the game to make it shine, and it shows. The level layouts are great and provide their own challenges. Even the stereotypes are great!

The game is a little short, but magic nonetheless. I can't wait for DLC (I can't believe my own ears saying this!) or for Huntdown II, the next chapter.

If there is one thing lacking in the game, it is the multiplayer. While the game supports local co-op, it does not have online multiplayer. This would not have been an issue in 1980, when all my

friends lived in a five mile radius, and it was awesome to have them over to play on my Atari 2600, whilst sharing peanut butter and apricot jam sarmies; today, most kids' friends are further away.

Simple mechanics, easy controls, great design and satisfying boss fights make this title easy to recommend. It is by no means perfect, but it has a LOT of potential to be a really awesome game. If ever you see a game advertise "retro arcade action" - they need to be compared to this.

Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

PATRONS

Monthly Patrons

Alex Crabtree
Alex Popescu
Andy Garay
Bill Berninghausen
Bob C
Brian Bogdan
CBinMV
Darren
Dennis Mack
Devin McPherson
Doug Bruce
Elizabeth K. Joseph
Eric Meddleton
Gary Campbell
George Smith
Henry D Mills
Hugo Sutherland
Jack
Joao Cantinho Lopes
John Andrews
John Malon
John Prigge
Jonathan Pienaar
JT
Kevin O'Brien
Lee Allen
Leo Paesen
Linda P
Mark Shuttleworth
Norman Phillips

Oscar Rivera
Paul Anderson
Paul Readovin
Rino Ragucci
Rob Fitzgerald
Roy Milner
Scott Mack
Sony Varghese
Taylor Conroy
Tom Bell
Tony
Vincent Jobard
Volker Bradley
William von Hagen
Jason D. Moss

Single Donations

2021:
Floyd Smith
Dale Reisfield
Jan Ågren
Linda Prinsen
melvyn smith
Frits van Leeuwen
Raymond Mccarthy
Robert Kaspar
Frank Dinger
Ken Maunder
Brian Kelly
János Horváth
Ronald Eike

John Porubek
Hans van Eekelen
Kees Moerman

HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users. Send your articles to: articles@fullcirclemagazine.org

We are always looking for new articles to include in Full Circle. For help and advice please see the **Official Full Circle Style Guide**: http://bit.ly/fcmwriting

Send your **comments** or Linux experiences to: letters@fullcirclemagazine.org
Hardware/software **reviews** should be sent to: reviews@fullcirclemagazine.org
Questions for Q&A should go to: questions@fullcirclemagazine.org
Desktop screens should be emailed to: misc@fullcirclemagazine.org
... or you can visit our **site** via: fullcirclemagazine.org

FCM#172
Deadline:
Sunday 08th Aug. 2021.
Release:
Friday 27th Aug. 2021.

For the Full Circle Weekly News:

You can keep up to date with the Weekly News using the RSS feed: http://fullcirclemagazine.org/feed/podcast

Or, if you're out and about, you can get the Weekly News via Stitcher Radio (Android/iOS/web):
 http://www.stitcher.com/s?fid=85347&refid=stpr

and via TuneIn at: http://tunein.com/radio/Full-Circle-Weekly-News-p855064/

Getting Full Circle Magazine:

EPUB Format - Most editions have a link to the epub file on that issue's download page. If you have any problems with the epub file, email: mobile@fullcirclemagazine.org

Issuu - You can read Full Circle online via Issuu: http://issuu.com/fullcirclemagazine. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu.

FCM PATREON : https://www.patreon.com/fullcirclemagazine