INSIDE RASPBERRY PI 4

The next generation Raspberry Pi 4 is a major advance over the Raspberry Pi 3 series. Many bottlenecks have been lifted, many features which the community requested introduced. It has the same well-known and -loved form-factor. But, you're going to need new accessories, and there are some other aspects of the new technology used you should be aware of.

In this article I will take a deep dive inside the Raspberry Pi 4, and share its exciting technical details with you. The views expressed in this article are entirely my own, backed up with careful research. Many thanks to the Raspberry Pi engineers & commercial team, which provided answers to my questions, and to many other approved resellers' questions, which allowed this article to become possible!

 And before you ask - the Pi 4 still available at 35 $!


A first look at the Raspberry Pi 4

THE HIGHLIGHTS & CORE FEATURES

The new Pi 4 is roughly three times faster than the predecessor model Pi 3B+, according to Eben Upton. It has up to 4 GB of RAM, two microHDMI ports, two USB 3.0 ports, and true Gigabit Ethernet. The Gbit Ethernet is engineered into the SoC, and does not share the bandwidth with the USB ports as on all previous models.

The Pi 4 is fully downward compatible, the GPIO header, PoE extension pins, camera & display ports are in the expected places and support the same kind of hardware as on the Pi 3B+.

You will need new accessories for the Pi 4:

- 3 A USB C power supply (the Foundation has its own - recommended)
- microHDMI to HDMI cable(s)
- new case (due to new connectors and different placement, Pi 3 cases will not fit)
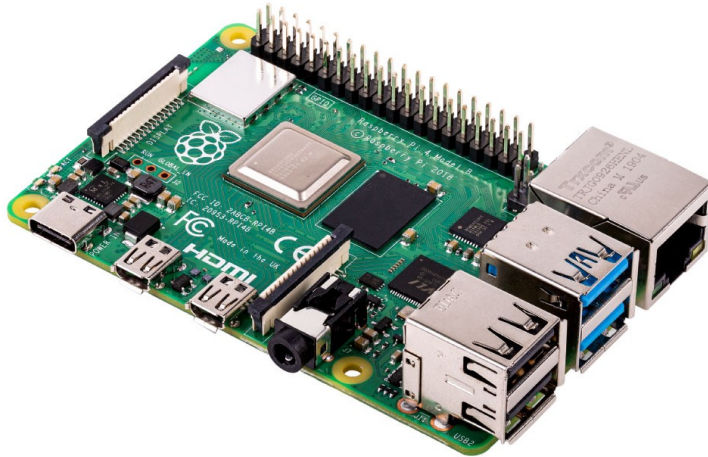
The picture shows our Comfort Set (Pi 4 / 4 GB RAM) with the new 3 A USB C power supply, two microHDMI / HDMI cables, and the new case. Order number: b-rpi4b.4g-psuc3a-ocasered-32gb-2xmicrohdmik-cat6

pi3g
www.pi3g.com

THE SPECS

There's so many new things to talk about - therefore let's have a look at the specs, as a structural foundation for our discussion of the new Pi 4 Model B. Afterwards I recommend you dive into the parts which interests you most, and skip the discussion of the parts which you're not interested in.

CPU / SoC	Broadcom BCM2711, quad-core Cortex-A72 (ARM v8) 64-bit SoC @ 1,5 GHz (28 nm)
Memory	LPDDR4, 1 GB / 2 GB / 4 GB variants. list prices are: • **1 GB - $ 35** • 2 GB - $ 45 • 4 GB - $ 55
Connectivity	2,4 GHz / 5 GHz IEEE 802.11 b/g/n/ac wireless LAN Bluetooth 5.0 / BLE Gigabit Ethernet 2 x USB 3.0 ports 2 x USB 2.0 ports
GPIO	Standard 40-pin GPIO header *(fully backwards-compatible with previous boards)*
Video & sound	2 x micro HDMI ports (up to 4Kp60 supported) 2-lane MIPI DSI display port 2-lane MIPI CSI camera port 4-pole stereo audio and composite video port
Multimedia	H.265 (4Kp60 decode) H.264 (1080p60 decode, 1080p30 encode) OpenGL ES 3.0 graphics
SD card support	microSD card slot for OS & data storage
Input power	5 V DC via USB-C connector (3A min.) 5 V DC via GPIO header (3A min.) Power over Ethernet (PoE)-enabled, requires separate PoE HAT
Environment	Operating temperature 0-50°C
Compliance	see https://www.raspberrypi.org/documentation/hardware/raspberrypi/conformity.md
Production lifetime	Pi 4 Model B will remain in production until at least January 2026

28 NM: SHRINKING INTO THE FUTURE

The new Raspberry Pi SoC is built on the 28 nm node. What does this mean for us?

So far, all previous Raspberry Pi SoCs were built on the **40 nm** process node. A "node" refers to a specific semiconductor manufacturing process and its design rules. The number used to indicate the length of a transistor gate manufactured using this process, but has lost this meaning on the smallest currently manufactured nodes due to marketing interests.

For a long time, reducing the size of transistors has been the secret behind the seemingly magical developments in computers and related IT technology: ever-increasing performance at sinking prices, and lower power usage. The cost per transistor used to drop with each smaller node, allowing the designer to either add many more transistors (and thus features), or sell the chips at significantly lower prices. Advances in the lithography (technology used to create these chips) allowed us to have microSD cards with 1 TB capacity, and powerful battery-powered multimedia computers in our pockets (smartphones).

Unfortunately, the technology needed to manufacture at smaller and smaller nodes significantly increases in complexity, and cost. A modern fab, where chips are manufactured, is a multi-billion $ investment. Currently only three companies are capable of manufacturing at the bleeding edge ("7 nm"): Samsung, TSMC, and Intel. These chips are used in high-end computers and mobile devices.

Why did this shrink not occur before?

Chips are developed for certain technology nodes. To change the node, a lot of the chip needs to be redesigned. Rather than giving us the same old a bit faster, the Raspberry Pi Foundation did the right thing, and completely overhauled the chip - even though we waited a bit longer for it, the new possibilities are worth it.

So why didn't the Raspberry Pi Foundation go all out and drop the node to the bleeding edge?

One simple reason: cost. The Raspberry Pi Foundation has always been about cost-efficiency. The sweet spot in semiconductor manufacturing is currently on the 28 nm node. Even though smaller nodes might enable the Pi to draw less power, run cooler, or have more performance, it would drive up the price - and the 35 $ price point is holy to the Raspberry Pi Foundation.

This also means that, compared to a smartphone (or competitor SBC) with the SoC manufactured on a smaller process node, the Pi 4 will use more energy for the same computations, and dissipate more heat.

CPU / SOC

WHY WE NOW GET MORE CPU POWER

The biggest advantage that the node shrink brings is to be able to put more complex chip designs (with more transistors) in the same physical space.

The second biggest advantage is to decrease heat dissipation required for the same compute power. A CPU design has to fit into a certain thermal design power envelope (TDP), **which is the maximum amount of heat it is designed to dissipate running real applications, without requiring additional cooling systems**.

The shrink to the 28 nm node, in addition to certain CPU architecture optimizations, therefore gave us a rough performance increase of three times. Still, there is a limit to what can be dissipated on a credit-card sized PCB, which is now synonymous with Raspberry Pi, and therefore a limit to what kind of maximum CPU power is possible.

Like the Pi 3B+ the Pi 4 dissipates the heat generated by the SoC in a very clever fashion by using a thick layer of copper in the PCB, designed to guide the heat and dissipate it across the entire Raspberry Pi PCB. For instance, you'll notice that the USB C power jack, the microHDMI ports, the LAN port and the USB ports feel quite hot to the touch during operation - whereas touching the SoC even feels painful. The BCM2711 SoC runs hotter than the SoCs on the previous Pis.

The SoC will monitor its own temperature (using built-in sensors) and throttle if it goes over certain temperature limits. In order to increase the performance, therefore, passive and/or active cooling measures can be implemented. For instance: a heatsink on the SoC, or an active fan (as it is implemented in the PoE HAT).

*The picture shows our **b-rpi4b.4g-hs** model - a Pi 4 B/4GB with a pre-mounted heatsink, available commercially.*

pi3g
www.pi3g.com

PERFORMANCE COMPARISONS TO PREVIOUS PI GENERATIONS

As on the previous multi-core Raspberry Pi models, the Pi 4 also has four cores.

When comparing CPUs it is very important not to look at just the clock rate, but also the design and efficiency of the CPU for different computing tasks. A rough way to compare and get a first feeling for the CPU are the Dhrystone Million Instructions per Second (DMIPs). The more, the better.

Also, DMIPs can be compared across platforms and CPU architectures (e.g. x86), while CPU frequency comparisons are meaningless across different architectures.

An important value to look at is DMIPs / MHz, which indicates how efficiently the CPU uses the clock cycles. A lot of the performance gain on the new Pi 4 is due to an increased DMIPs/ MHz value.

The following values are calculated values, using numbers from Wikipedia as a reference.

Raspberry Pi Model	Pi Zero W	Pi 2 (first version)	Pi 3B+	Pi 4
SoC	BCM2835	BCM2836	BCM2837B0	BCM2711
ARM core(s)	1 x ARM1176JZF-S	4 x Cortex-A7	4 x Cortex-A53	4 x Cortex-A72
Architecture	32bit (ARMv6Z)	32bit (ARMv7-A)	64bit (ARMv8-A)	64bit (ARMv8-A)
Clock rate	1 GHz	0.9 GHz	1.4 GHz	1.5 GHz
DMIPs / MHz	1,25	1,90	2,24	4,72
total calculated DMIPs	1.250	6.840	12.544	28.320
performance compared to Pi Zero	1	~ 5,5	~ 10	~ 22,6

Ref:

- https://en.wikipedia.org/wiki/List_of_ARM_microarchitectures
- https://en.wikipedia.org/wiki/Comparison_of_ARMv8-A_cores

As the CPUs are optimized for real-world usage scenarios, like web browsing, and not synthetic benchmarks, actual "felt" performance will be even higher - web browsing, for instance, feels notably smoother. Other factors, like the RAM being much faster (see below) also contribute to a higher execution speed - the CPU can work faster, as data becomes ready (= transferred to the internal CPU cache) earlier.

One important new feature in the Pi 4's ARM core powering the higher DMIPs / MHz is the **out-of-order execution** ("*pipelined processor with deeply out of order, speculative issue 3-way superscalar execution pipeline*"). This allows the processor to continue executing other code in the application, while waiting for required data to become ready, or to execute a different kind of code in parallel while certain CPU subsystems are busy. The instruction cycles are not wasted, the CPU is thus faster.

The new CPU is a significant upgrade and plays in a different league. Looking just at the DMIPS it roughly corresponds to an Intel Core 2 Extreme X6800 (2-core) CPU, introduced in 2006 (27.079 DMIPS at 2,93 GHz).

My still quite powerful 2012 notebook, on which I am writing this, has an Intel Core i7-3632QM CPU (22 nm), for a calculated DMIPS value of 99.750 - roughly **only three times more** than the Pi 4.

pi3g
www.pi3g.com

BOOTUP NEWS: ONBOARD EEPROM

The Pi 4 has an onboard 4 Mbit (=512 KByte) EEPROM, which currently contains the second stage bootloader. This bootloader can be updated for future boot scenarios. At launch it is likely that the Pi 4 will not be able to boot over USB 3.0 / GBit Ethernet, etc. This EEPROM will ensure that Pi owners will benefit from changes in the boot code.

CPU / SOC Q & A

Why not more cores?

More cores need more die space, require more power and the additional heat needs to be dissipated. Especially the heat dissipation is a challenging aspect of the design. Thanks to the shrink to the 28nm node, and the more efficient CPU, the Pi 4 can perform two to five times faster without using and dissipating two to five times the energy. As the cores which are on-board will start to throttle due to excess heat, adding more cores would exacerbate the problem - and not really provide additional, sustained performance.

What about the new SoC number?

Indeed, the SoC number BCM2711 for the Pi 4 is "lower" than on the 3B+ (BCM2837B0) or the Zero W (BCM2835). This can be explained by 2708 (BCM2708) being the product family of the first Pi generation, and BCM2835 a specific implementation (package). The Pi 2 was BCM2709 (name of the silicon die) and BCM2836 the package, and the Pi 3 BCM2710 / BCM2837. Therefore actually BCM2711 (Pi 4) is a step forwards, the next generation after Pi 3 (BCM2710).

Ref:

- https://raspberrypi.stackexchange.com/questions/840/why-is-the-cpu-sometimes-referred-to-as-bcm2708-sometimes-bcm2835

How does using the new Pi actually feel like?

It feels so smooth (especially on the 4 GB version) that you tend to forget that you are using a single board computer. Web surfing feels just like on a Windows desktop on your big desktop or notebook computer, with minor UI differences. Using LibreOffice feels natural. YouTube, Facebook, mail - everything just works as expected. The Pi will be great as a desktop for many users! (Go for the 4 GB RAM version).

Will we switch to a 64 bit Raspbian?

It will not happen right now, maybe sometime down the line. There are diverse software and deployment / compatibility tasks to be taken care of, there are some advantages and some disadvantages (e.g. program code requires more space!). At the moment most users would probably not see a tangible benefit from.

A main advantage of 64 bit application space is the ability to handle huge files, and/or a lot of files for applications. E.g. graphics artists, 3D rendering, video editors, etc. These are quite specialized applications which the Pi 4 is not geared towards.

Why is the Pi 4 appearing now, we were thinking 2020?

This particular SoC became ready earlier than expected, thus the Foundation launched as soon as they could.

MEMORY
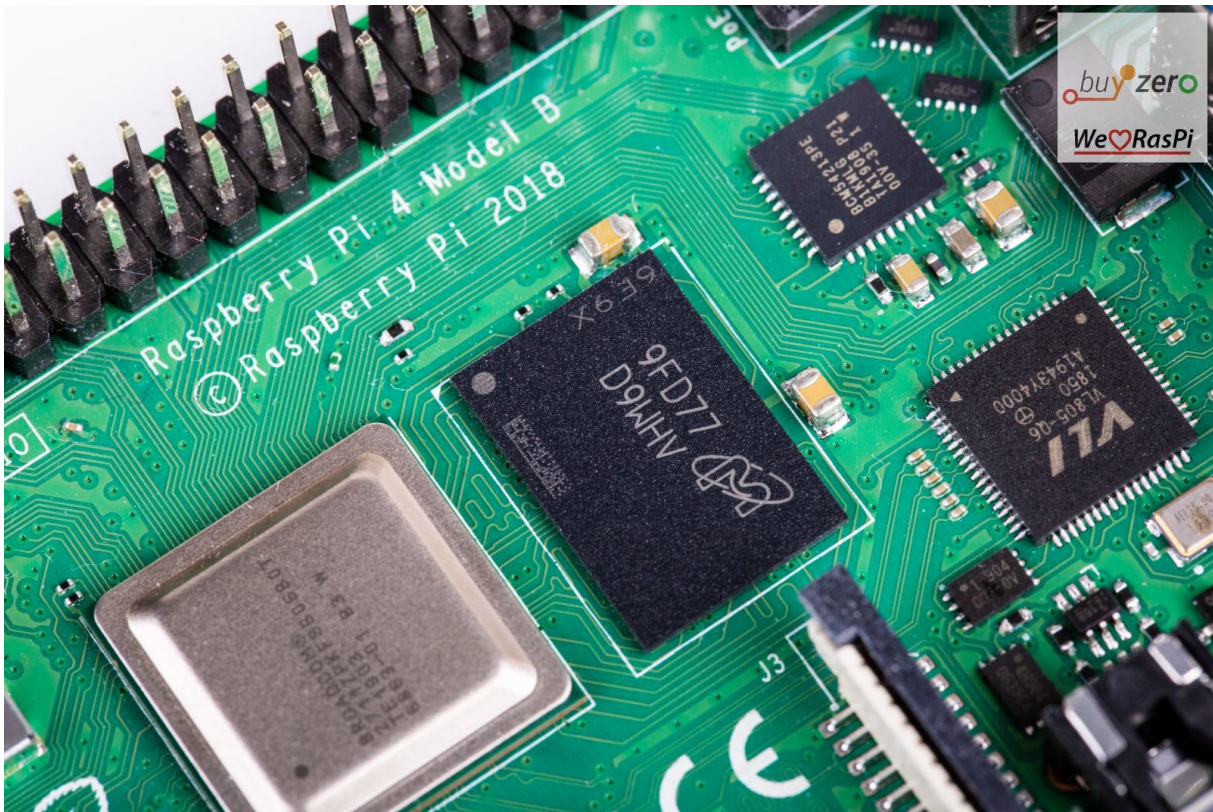
The Raspberry Pi 4 Model B is the first Pi to come in **three memory size flavors**:

- Pi 4 / 1 GB
- Pi 4 / 2 GB
- Pi 4 / 4 GB

The 1 GB variant has the lowest price, while the 4 GB variant offers the best performance. For memory-hungry applications, such as web browsing, 4 GB is recommended. For embedded applications, such as digital signage, or headless usage the 1 GB or 2 GB model should be good enough.

The memory technology has been upgraded to LPDDR4 (Pi 3B+ and previous models use LPDDR2). The SoC was designed to have the memory on the top side of the Pi, next to the SoC. The memory on my sample board was supplied by Micron - it is the chip with the M and the swirl around it. The second code on the chip, D9WHV, can be entered into this online-tool for decoding:

https://www.micron.com/support/tools-and-utilities/fbga

The part number shows up as MT53D1024M32D4DT-053 WT:D

Please note that the density for this part is specified as 32Gb. This is actually in Gbit - to convert to Gbyte, divide by 8 (= 4).

The LP-DDR4 standard was published by JEDEC on the 25th of August 2014. It introduced significant improvements compared to LP-DDR2 on previous Pi generations.

The Raspberry Pi Foundation went from LPDDR2-900 (0,9 GT/s) to LPDDR4-2400 (2,4 GT/s), which is a ~ 2.5 increase in theoretical bandwidth. According to Eben Upton the new IP (SoC design) is pretty good at getting all that bandwidth out.

MEMORY Q & A

Why not supply all Pi 4 with 4 GB of memory?

More memory costs more money. That is why we get a choice of 1, 2 and 4 GB, according to our budgets and actual memory requirements. The $ 35 price point could not have been met with a 4 GB Raspberry Pi.

What are these GT/s?

LPDDR stands for Low Power Double Data Rate. Data is sent twice per clock cycle, thus allowing for lower clock frequencies. The designation LPDDR2-900 for instance means 0,9 GT/s, the actual data line clock runs at 450 MHz. The GT/s are Gigatransfers per second - in this case we have 450 million rising clock edges per second and 450 million falling clock edges per second (for the 450 MHz signal), adding together to 900 MT/s (megatransfers per second) or 0,9 GT/s. The important thing to know is that these GT/s do not specify the number of bits carried in each operation (as the memory bus carries several bits). You have to multiply the number of bits to get the actual bandwidth in GB/s which is possible for the technology.

Ref: https://en.wikipedia.org/wiki/Transfer_(computing)

What RAM size should I go for?

It will depend on your application. For web surfing and other desktop usage, the 4 GB version is highly recommended. Web browsers use a lot of memory, especially for multiple tabs open simultaneously.

The 1 GB version will be, due to lower cost, the ideal choice for embedded and industrial applications (headless / digital signage).

The 2 GB version might be OK for low-cost dual-screen thin client deployments.

When in doubt, and if you have the budget, go for the higher memory.

Will the LPDDR4 memory require a heatsink?

No, it does not require a heatsink when used within normal operating conditions.

CONNECTIVITY

The Pi 4 has two features which the community desperately wanted for a long time now: true Gigabit Ethernet, and 2 x USB 3.0.
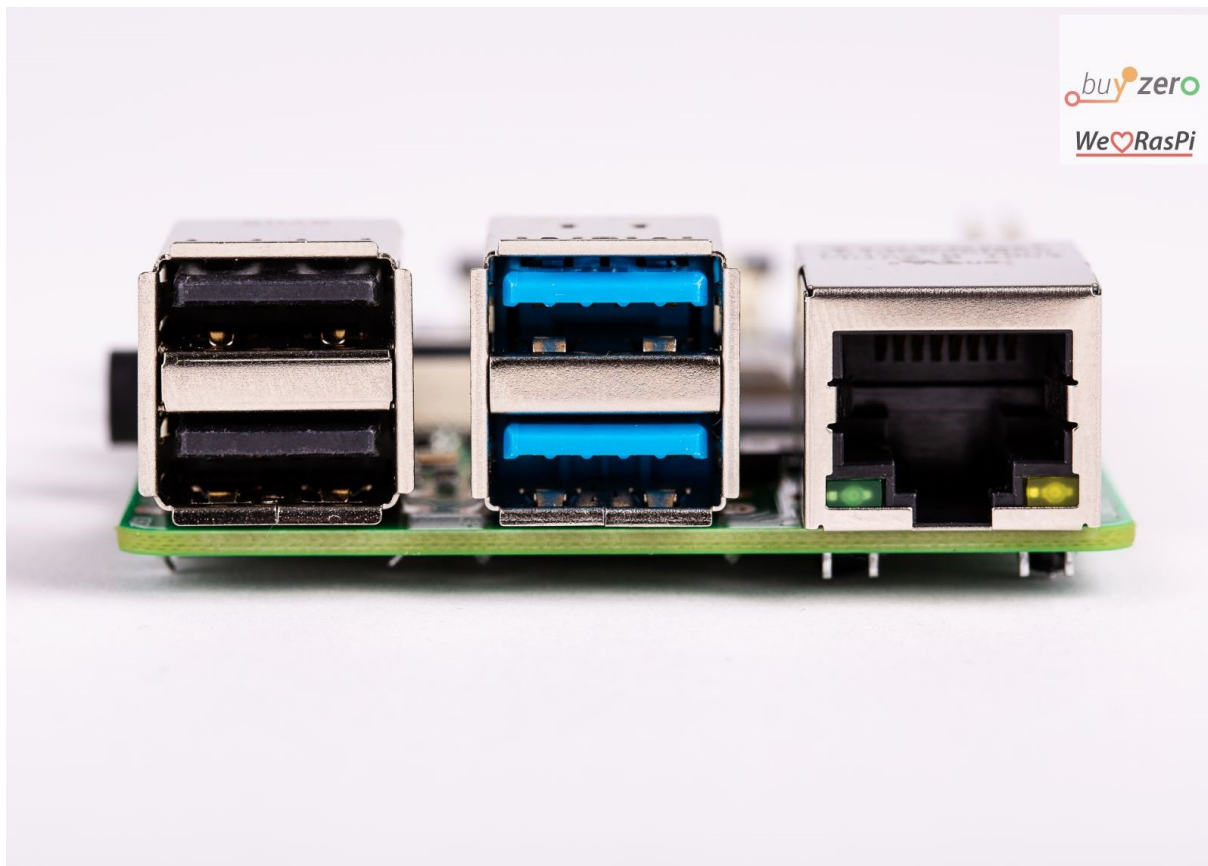
ETHERNET

Yes, this is true Gigabit Ethernet, connected to the SoC - not through USB! The LAN port has moved to be on the side of the GPIO and PoE pins: one of the reasons you will need a new case for the Pi 4.

Finally, the Ethernet's and the USB's bandwidth does not have to be shared anymore - this is wonderful news for NAS-builders, backup solutions, self-built cloud storage, network video streaming and much, much more.

The Ethernet will of course continue to provide power over Ethernet / PoE functionality (with the additional PoE HAT), and support Ethernet boot (potentially to come only with a later software update).

The Pi 3B+ was limited to 330 Mbit/s Ethernet due to its internal USB 2.0 bottleneck.

The Pi 4 thus has a roughly **three times** increased Ethernet bandwidth (1 Gbit/s = 1000 Mbit/s). Another advantage is that the Ethernet packets do not need to be repacked into USB packets anymore, thus latency and processing overhead are decreased on the Pi 4. The new Raspberry Pi product thus fully caught up with modern Ethernet hardware and networks, and is therein state-of-the-art.

The next step up, 10 GBit/s Ethernet is still (as of writing this) enterprise / server class hardware.

More juicy technical details on the Gbit Ethernet:
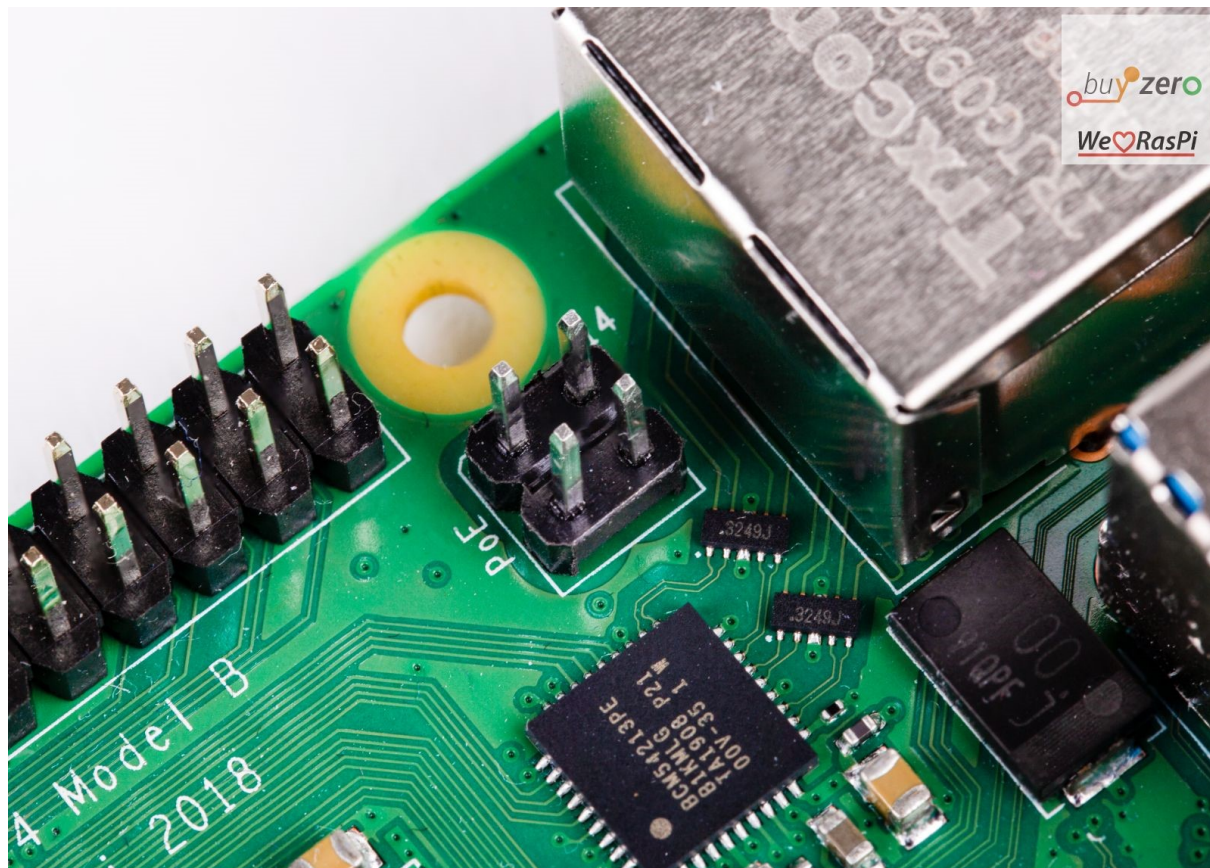
Technically speaking, the Ethernet jack is connected to a part from Broadcom, which carries the inscription **BCM54213PE**.

This BCM54213PE is an Ethernet transceiver (PHY), which performs all the necessary analog & physical functions for the Ethernet. The data is passed to and from the BCM2711 SoC and its built-in MAC (media access controller) using an interface called RGMII. The BCM54213PE thus gets a stream of bytes from the SoC's MAC and converts it into analogue signals on the wires, and vice-versa.

RGMII stands for Reduced Gigabit Media-Independent Interface. The reduced is not referring to reduced speed, rather to reduced amount of data lines (12 instead of 24) necessary to connect the two chips together. It is a standardized interface, so MACs can speak to different PHYs, e.g. for fiber optic media etc., without having to redesign the MAC.

Ref: https://www.broadcom.com/products/ethernet-connectivity/copper-phy/gigabitphy/bcm54213pe

Some words on the throughput of GBit Ethernet:

There are different data frame overheads to be deducted from the full theoretical capacity of the GBit Ethernet. For the default maximum frame size of 1518 bytes for Ethernet (with an MTU of 1500), they are as follows:

- Interframe Gap (12 bytes)
- Preamble (8 bytes)

- Ethernet header (14 bytes)
- IP header (20 bytes for IPv4)
- TCP header (20 bytes)
- Frame check sequence: CRC32 checksum (4 bytes)

a total of 78 extra bytes for each 1460 bytes of data. Therefore the practical maximum throughput is around **118 MB/s**, lower than the specified bandwidth of 125 MB/s, an efficiency of around 94 %. A higher efficiency of up to 99 % can be achieved using so-called Jumbo Frames, by increasing the MTU to 9000 - transporting 8960 bytes of data per packet.

Ref: https://www.gigabit-wireless.com/gigabit-wireless/actual-maximum-throughput-gigabit-ethernet/

USB 3.0

USB 3.0 is not a direct connection to the SoC - there's a chip in between.
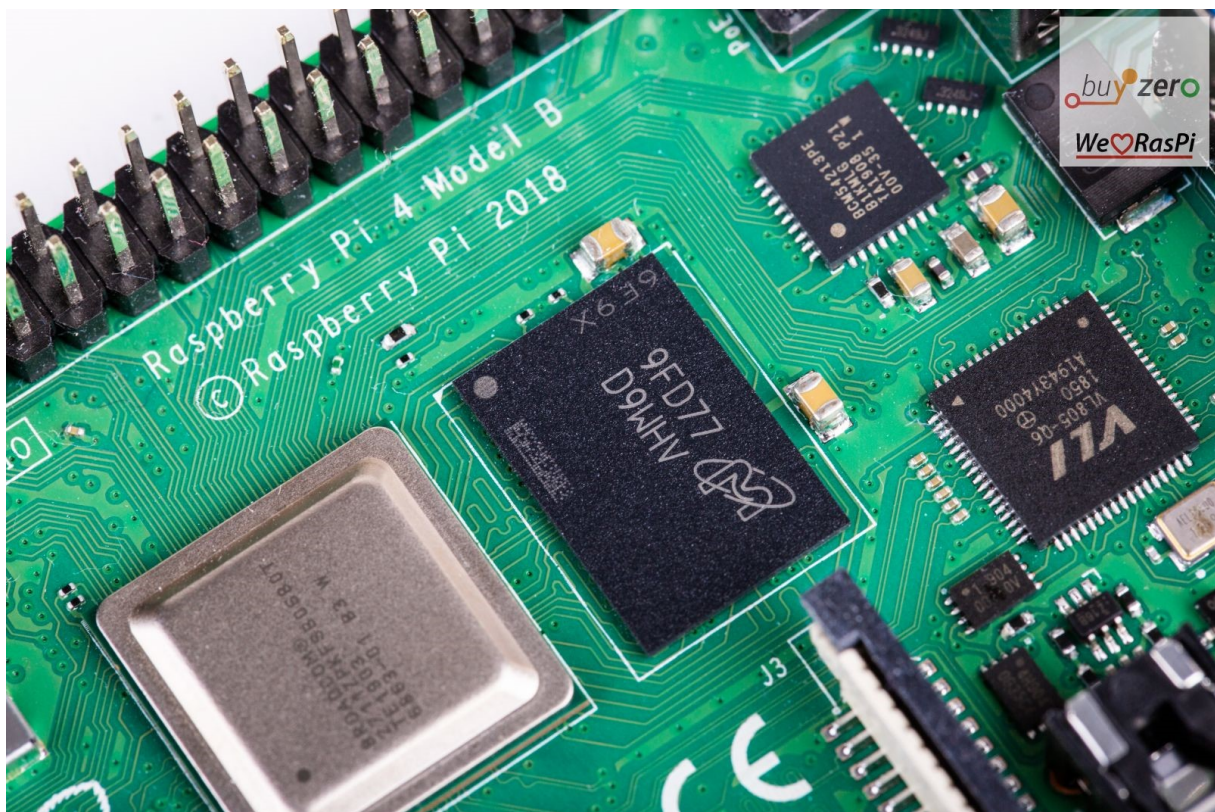
The chip is **VLI VL805-Q6** from Via Labs (VL805).

This is a PCI Express to 4-port USB 3.0 super speed host controller. Upstream, it has a PCI Express 2.0 x1 (one lane) interface.

PCI Express 2.0 is the new fast communication port on the Raspberry Pi 4's SoC, the BCM2711: basically the equivalent of the USB 2.0 port on the previous Pi generations. PCI Express is used on most modern computers as the internal fast interface of choice, albeit not necessary in version 2.0. The BCM2711 supports one PCI Express 2.0 lane.

USB 2.0 provides a bandwidth of 480 Mbit/s or 60 MB/s (Megabytes/s). This is actually the total bus bandwidth, not the maximum real bandwidth available for transferring data after subtracting bus overhead.

The maximum real bandwidth on the USB 2.0 bus is estimated to be about **53 MB/s**. With USB 2.0, the bandwidth is shared between the up and down direction (half-duplex). Data only flows in **one direction** at a given time.

PCI Express 2.0 supports 5 GBit/s bus bandwidth per lane. It uses 8/10bit coding to ensure reliable transmission of the data, which allows for **500 MB/s** of actual data transfer rate per lane. It is important to note, however, that this bandwidth is provided simultaneously and independently in both directions (up / down).

Therefore the maximum throughput has increased about **ten to twenty-fold**, depending on whether simultaneous up/down transfers are being considered. Gone is the USB 2.0 bottleneck of previous Pi's.

So what about USB 3.0 "SuperSpeed" (aka USB 3.1 Gen 1)?

USB 3.0 has a data transfer rate of up to 5 Gbit/s (625 MB/s) (total bus bandwidth). The theoretical real bandwidth after subtracting overhead, which is available for data transfer is 4 Gbit/s (**500 MB/s**). As PCI Express, USB 3.0 is a full-duplex standard. Data can flow in both directions at the maximum data rates simultaneously.

Therefore **one** USB 3.0 port will be able to saturate the single PCI Express link perfectly in both directions (all the bandwidth available for real data transfer can be carried on the PCI Express, as the bus coding overhead is unpacked by the VLI controller).

Don't expect to be able to use both USB 3.0 ports full throttle at the same time, though.

Ref:

- https://www.via-labs.com/product_show.php?id=48
- https://www.cypress.com/file/134171/download
- https://en.wikipedia.org/wiki/List_of_interface_bit_rates

What about PCI Express 3.0, 4.0, 5.0?

PCI Express 4.0 was introduced 2017, and PCI Express 5.0 in 2019. They simply were not standarized yet, when work on the Pi 4's SoC most likely started. PCI Express 3.0 must have been available by then (it was introduced in 2010); probably PCI Express 2.0 was chosen for cost / performance tradeoff reasons; possibly PCI Express 3.0 compatible USB 3 controllers would be more expensive.

I assume that hereby we have an upgrade path for future Raspberry Pi generations, once component cost comes down.

Why not USB 3.1 Gen 2 / USB 3.2?

The one PCI Express 2.0 lane would not support the required data rates.

Why no USB C connectors?

There are still more devices using "traditional" USB plugs, for instance your average office mouse and keyboard, USB stick, etc. USB-C connectors are more expensive.

The power connector is already a USB-C connector, this way mix-ups are not just unlikely, but not possible (only one possibility to connect the power supply for the Pi).

Why not four USB 3.0 ports, but just two?

Only two of the theoretically possible four USB 3.0 ports with the new PCI Express to USB 3.0 bridge are used, the other two being downgraded to USB 2.0, mainly because of routing constraints. That means, that it is difficult to squeeze so many PCB traces on the Raspberry Pi and its PCB layers, because of all the functionality which is built-in. USB 2.0 was chosen because it requires less traces.

It is not a major loss, though, if you consider the typical USB usage scenario on a Raspberry Pi: you will have to attach mouse and keyboard, and these do not require USB 3.0 ports. Also, consider that you will not be able to saturate the bandwidth of the two existing USB 3.0 ports, as they share one PCI Express 2.0 lane back to the SoC.

What kind of power consumption is possible on the USB ports?

The maximal downstream current consumption is 1.2 A, limited by the Pi. It is shared across all four ports - you can therefore draw a maximum of 1.2 A for a single port if no other USB devices are attached. If you need more power from the USB ports, consider using a powered USB hub.

Another alternative is the Pi Zero - it does not have polyfuses and uses much less power itself. The attached USB device can draw nearly as much power as you are able to supply to the Pi Zero.

Are the USB 3 ports and the Gigabit Ethernet really independent?

Yes, indeed they are. You can saturate the Gigabit Ethernet and the PCI Express interface simultaneously. As discussed, USB 3.0 bandwidth is shared across the two USB 3.0 ports.

This should enable NAS solutions.

What kind of real-life performance can I expect from the USB 3.0 port?

Reading from SSD has been tested at over 350 MB/s. Your mileage may vary, depending on the specific SSD and read patterns you are using.

Note that the actual data throughput on the GBit Ethernet port after accounting for overheads is about 118 MB/s (see the Ethernet section). Therefore you will not be able to push the data at the speed at which you can read it from the SSD over the GBit network. For normal real-life small office or home scenarios, this should still be plenty. Working on NAS filesystems through Gbit Ethernet feels mostly like working on files locally.

Here's a little preview what you can expect for copying between two drives, both attached via USB 3.0 to the Raspberry Pi 4 (for 4.4 GB of data).
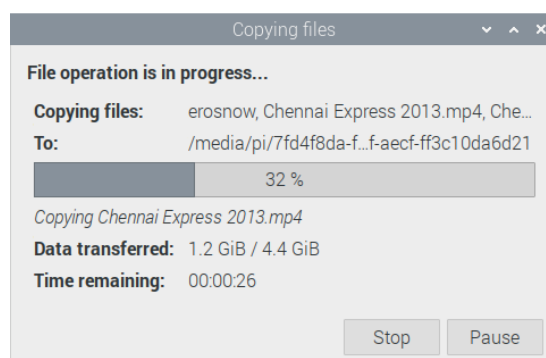
What is this thing with the Megabits and Megabytes?

A byte contains 8 bits. Historically this is due to 8 bits being required / specified to encode a single character (e.g. "A"). We have to be careful which of the ones we talk about in throughput measurements; else we end up comparing apples to oranges.

Is it possible to get USB OTG on the USB 3 / USB 2 ports?

No, the VL805 is host-only. If you need USB OTG functionality, the legacy USB 2.0 port on the Pi 4 has been routed to the USB-C connector.

A SIDENOTE ABOUT SIDELINED SATA

Unfortunately, the Pi 4 does not provide native SATA connectivity. To attach SATA hard drives to the Pi 4, you can use a SATA to USB 3.0 bridge. Please note, that the 3 A power supply will not be able to power the Pi and your external 2.5'' SATA disk (e.g. SATA SSD) in some cases. I suggest to use 3.5'' enclosures and their own power supply, or an external powered USB 3.0 hub.

The picture shows one of our NAS sets, order code b-rpi4b.4g-hs-psuc3a-ocasered-32gb-cat6-satacase-hdd4tb

Why was SATA not included?

Each additional component on the Pi increases the cost. Even a SATA connector costs money, and as not every person using a Pi wishes to connect a SATA hard drive to it, it has to stand back in place until more "important" features (relevant to more people) fit within the $ 35 budget. Possibly also the new USB 3.0 interface, and the easy extensibility it provides were a consideration. Furthermore, USB 3.0 is an external interface, whereas SATA

is an internal interface. Normal Raspberry Pi cases will allow it to be used as a small computer to which you can attach external drives. SATA drives require special mounting enclosures, etc.

WLAN

The WLAN chip (Cypress CYW43455) lives in a metal housing (it helps the modular certification), this time it is not embossed with a Raspberry Pi logo. It is the same chip as on the Pi 3B+! It is connected to the SoC using SDIO, as on the Pi 3B+. It continues to use the Proant AB PCB antenna. No changes here.

- 2,4 GHz / 5 GHz
- IEEE 802.11 b/g/n/ac

The WiFi should have a similar performance to the Pi 3B+, also according to various throughput tests.

Have a look at the following page for performance tests of the Pi 3B+:

https://www.phoronix.com/scan.php?page=article&item=raspberry-pi3bplus-wifi&num=2

- UDP payload - about 50 Mbit/s on 2,4 GHz and about 100 Mbit/s on 5 GHz should be reachable
- TCP payload - about 50 Mbit/s on 2,4 GHz and about 75 Mbit/s on 5 GHz should be reachable

For optimal performance (throughput, latency, and less disturbances by neighbors in densely populated areas) we recommend to use the Pi 4 with a 5 GHz network, if one is available.

BLUETOOTH 5.0

There are over 8 billion Bluetooth products on the planet, and chances are you might want to connect one of them to your Pi. In the Pi 4, Bluetooth has been upgraded to Bluetooth 5.0 - or Bluetooth 5 in marketing speak.

Bluetooth 5 was announced on 16th of June 2016, focused on Internet of Things. Bluetooth low energy was introduced with Bluetooth 4.0, it is intended for embedded devices which want to send a small amount of data and go to sleep. The main focus is in the name: low power usage. Bluetooth low energy was developed to be interoperable with the "main" Bluetooth, which in Bluetooth 3.0 + HS had data rates of up to 24 Mbps. The new features of Bluetooth 5 are specifically designed for Bluetooth with low energy functionality, therefore it must be compared with Bluetooth low energy 4.2.

Bluetooth 5.0 low energy is two times faster (2 Mbit/s burst ~ 1,7x faster real-life) and has a four times bigger range (200+ m) than Bluetooth 4.2 (available on the Pi 3B+). It is an either-or, so devices have to choose the higher speed, or the bigger range. The faster data rates could enable, for instance, faster Firmware Updates on connected Bluetooth IoT sensors. Also, with Bluetooth 5 there is enough bandwidth for audio to be streamed over Bluetooth Low Energy. The bigger range will possibly enable you to use devices beyond your flat, e.g. sensors in the garden or garage. Devices should be able to communicate with up to 125 Kbps at the extended range. Of course, your sensors would also have to support the new standard.

The "0.9" mesh networking standard, part of Bluetooth 5, allows the creation of large-scale device networks. Data is passed along from sensor to sensor, until it reaches a gateway (like the Pi 4). Up to 32.000 nodes are supported, theoretically. This competes directly with other mesh-networking standards, such as Thread, ZigBee and Z-Wave. This is a major point for smart-home applications, based on Bluetooth mesh-networking compatible sensors, smart-building, industrial factory automation and many other sensor-based applications. A major advantage with competing standards is that this technology is already built-in in many devices which support Bluetooth 5, you do not need any additional hardware radio adapters. Sensors participating in the

mesh network only need to wake up once every four days, or if they have data to transmit - thus the power consumption stays low, the nodes can operate for many years off a coin cell.
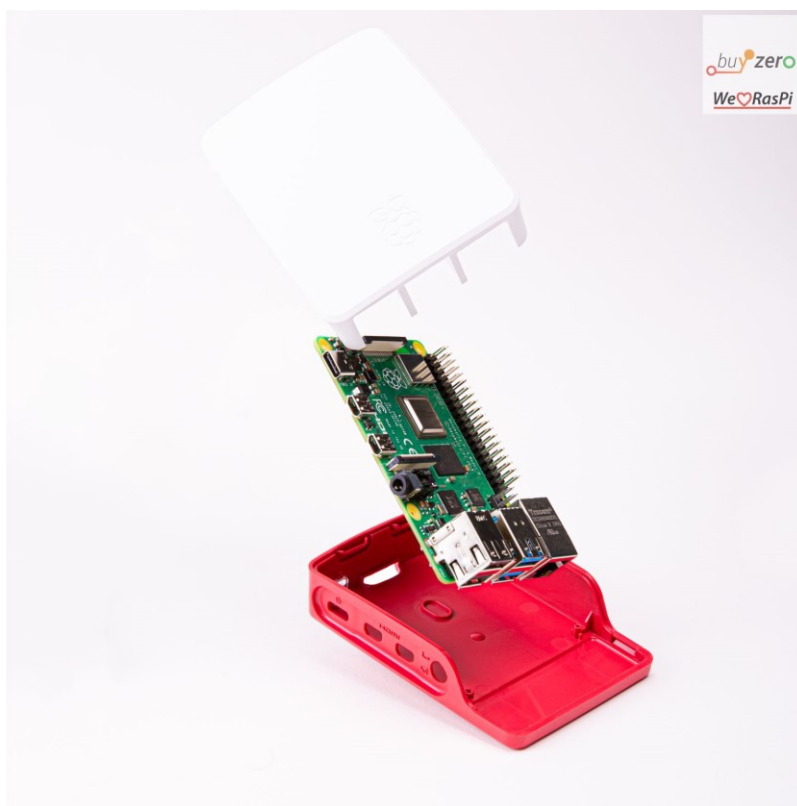
Bluetooth 5 increases the data broadcasting capacity of transmissions eightfold by increasing the advertising data packet lengths to up to 255 bytes per packet ("advertising extensions"). Previously this had been capped at 31 bytes. This allows it to use the available broadcasting channels on the 2.4 GHz band more efficiently, specifically for Beacons which leverage this Bluetooth Low Energy state. Beacons stay in the advertising state, broadcasting data for other devices to read. The extension to 255 bytes allows to transmit more data, introducing new possible applications.

Last but not least, it adds some more location aware features, which help with indoor navigation / proximity services, asset tracking services, etc. This, together with the increased broadcasting capacity should be interesting for embedded developers, building beacon technology on top of Raspberry Pi.

Ref:

- https://en.wikipedia.org/wiki/Bluetooth#Bluetooth_5
- https://www.techradar.com/news/networking/bluetooth-5-everything-you-need-to-know-1323060
- https://www.androidauthority.com/bluetooth-5-speed-range-762369/
- https://www.mouser.com/pdfdocs/bluetooth-5-faq.pdf
- https://www.bluetooth.com/blog/introducing-bluetooth-mesh-networking/
- https://www.bluetooth.com/bluetooth-technology/topology-options/le-mesh/mesh-faq/
- https://www.paratuspeople.com/news/bluetooth-5-and-its-role-in-iot/38779/

GPIO HEADER

The GPIO expansion port powers a diverse range of accessories for the Raspberry Pi, making it into the tool of choice for factory automation, makers, educators and others. Hobbyists and professionals alike. The GPIO port allows you to turn a generic Raspberry Pi into a custom machine, for instance to communicate using the CAN bus (with the PiCan 2 series), to show blinking lights (blinkt! by Pimoroni), to deliver high quality audio, and much more. It has been an essential core feature of the Pi, and is included in all modern Pi's, starting with the Pi 1B+, in its 40 pin glory.

GPIO connectivity still works at 3V3 - and all functions are available which have been available on the previous GPIO port. As a side note, the Broadcom SoC provides native 3,3 V IO. Since there is no adapter chip in between the GPIO port and the SoC you should expect the same kind of timing behavior and capabilities.

The GPIO header has always been "muxable", that is, different functions could be selected on each individual GPIO pin by software. The previous generation Broadcom SoCs (Pi 1, 2, 3) were consistent in the extra functionality they exposed (SPI, I2C, PWM, etc.) - these could be selected by the alternative ALT0 - ALT5 pin functions on the GPIO pins.

The new BCM2711 adds some new interface functions on ALT6 and ALT7:

- two more I2C ports
- four new independent UARTs (all PL011s)
- PWMs can be used in parallel with audio playback, the audio circuit now has its own PWMs

This is not an exhaustive list, additional features will be documented in time.

GPIO Q & A

Is I3C supported?

No.

Is USB over GPIO supported?

No.

Is the miniUART still coupled to the clock, and therefore susceptible to frequency scaling?

Yes, it is - for compatibility reasons, the behavior was not changed. However, there are new UARTs available which are decoupled.

Has the I2C clock stretching bug been fixed?

Apparently yes, it has been fixed.

Is there more information about the new UARTs?

There are four new UARTs, PL011. They can be muxed on the GPIOs 0-15 as follows:

- GPIO 0-3 = UART2 TXD, RXD, CTS, RTS
- GPIO 4-7 = UART3 TXD, RXD, CTS, RTS
- GPIO 8-11 = UART4 TXD, RXD, CTS, RTS
- GPIO 12-15 = UART5 TXD, RXD, CTS, RTS

VIDEO & SOUND

The Raspberry Pi 4 can drive two displays, and supports 4K resolution. 4Kp60 is possible, but only on one display. 4Kp30 might be supported on both displays simultaneously in the future, 1080p60 definitely is. The reason for this limitation is the sheer volume of data which needs to be pushed around.

There are two important restrictions:

- Pixel clock must be < 600 MHz
- overall framebuffer size (for all displays) must be < 7680 x 7680 - therefore dual 4096 x 2160 displays won't work

A 4096x2160 with a pixel clock of 594 MHz should therefore work in theory. Most consumer displays, which are referred to as 4 K displays, have a resolution of 3840x2160 pixels.

In order to drive two displays, the connectors had to shrink - gone is the HDMI connector, two microHDMI connectors are in its place. The same 4-pole stereo audio and composite video port is still on-board. The Foundation is keeping the board backwards compatible this way, and allows people with very old screens (TVs from the Stone Age - think poor countries, where the tech budget is limited!) to still drive them.

The ports are compatible with HDMI v2.0, a step up from the HDMI v1.3 on the previous Pis. HDMI 2.0 increases the maximum bandwidth to 18.0 Gbit/s. This allows it to carry 4K video at 60 Hz with 24 bit color depth. Not all HDMI cables are capable of supporting the full bandwidth - in 4K applications we therefore to recommend to use high quality HDMI cables, like the one from the Raspberry Pi Foundation.

As mentioned above, the audio circuit has now its own dedicated PWMs. You can therefore use the PWM on the GPIO at the same time. There still is no dedicated audio hardware chip, the Raspberry Pi Foundation is very happy with the audio quality as it is after the work on the driver. The audio has its own isolated 3v3 power rail.

The volume level is good for comfortably listening on headphones without additional amplification; the subjective quality perceived by me is good as well. For even better quality with a dedicated audio card, various solutions are available on the market (e.g. the pHAT DAC by Pimoroni).
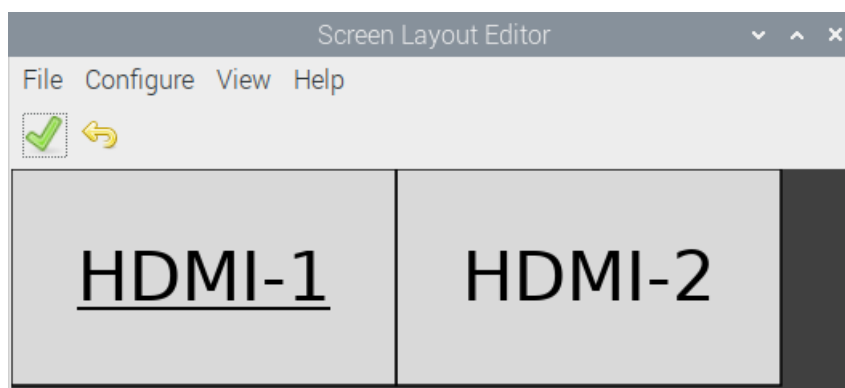
Q & A

Any changes to the CSI / DSI ports?

No changes known so far.

How will the two displays be managed?

There's a simple tool in Preferences -> Screen Configuration which allows you to set up how the screens are connected physically (i.e. where the boundary between the screens is and where the mouse should move).

Can I continue to use my HDMI cable?

To use your HDMI cable you would need a microHDMI to HDMI adapter. Two of these will not fit next to each other on the microHDMI ports of the Pi 4. Furthermore, it will put significant mechanical stress on the microHDMI port. To add on top of this, as discussed, legacy / cheap HDMI cables might not be able to carry the necessary bandwidth for 4K 60p video. We therefore highly recommend to get high quality microHDMI to HDMI cables for your new Raspberry Pi 4.

Is Ethernet over HDMI supported (HDMI Ethernet channel / HEC)?

No, it is not supported.

Is it possible to combine the DSI (7'' display), composite and HDMI?

composite and HDMI cannot be combined (as of now).

DSI (the official 7'' display) and HDMI-0 can be used together.

MULTIMEDIA

The secret to decoding video on SoCs without overtaxing the CPU or going outside the thermal budget is to have dedicated, specialized hardware blocks for video decoding.

The Pi 4 platform adds support for the modern H.265 codec in hardware, for up to 4Kp60 decode.

There are some limitations: as mentioned above, the 4Kp60 is possible for **one** video stream. The video must be encoded in H.265 - even though H.264 theoretically supports 4K resolutions, the Pi's hardware does not support them on H.264. Also, the new H.265 hardware block is for decoding only - to encode video, for instance from the Raspberry Pi camera, you will continue to use H.264, at 1080p30 max resolution.

VIDEOCORE VI

The Pi 4 has seen an upgrade to VideoCore VI.

On the SoC, the VideoCore V (VC5) 3D block is used (in an extended form). VideoCore V is used in many Broadcom set-top box chips. An evolved version, which also supports OpenGL ES 3.2 and Vulkan features was called VideoCore VI; the previous Pi's, with the VideoCore IV supported only OpenGL ES 1.1 and 2.0.

The hardware video scaler (HVS) used in the Pi 4 is a VideoCore V era component, so there is a mix of V and VI in the SoC.

The MPEG-2 and VC-1 Hardware decoders have been removed (or disabled), so keys for them are not going to be available. The respective codecs can be easily decoded on the ARM in software.

The H.265 / HEVC decoder is a HEVCv2 Main 4:4:4 10 design supporting bitstreams up to profile 5.1. Maximum picture size is 4096 x 4096.

The VideoCore 6 is about 4 times faster than the Video Core 4.

The VideoCore is one of the unique points about the Raspberry Pi and its software ecosystem. Most other SoC manufacturers use Mali, the GPU IP provided by ARM. Acceleration work and software drivers written for the VideoCore will thus benefit the Raspberry Pi platform, but not the competitors. Good software support has always been a strong point for Raspberry Pi, and VideoCore allows them to keep this advantage without supporting freeloaders.

MULTIMEDIA Q & A

What is OpenGL ES, what is Vulkan, and why does it matter?

Vulkan is a low-overhead, cross-platform 3D graphics and computing API. It was designed by the same group (Khronos Group) which developed OpenGL ES, as a kind of technology successor to it and OpenGL.

The important bit here being the cross-platform compatibility. Developers write games against APIs (application programming interfaces). These specific APIs are important for fast, hardware-accelerated rendering of 3D scenes, without having to go through normal operating system layers as normal applications like text editors, and terminal emulators do.

If each platform, each operating system, and each graphics card would have its own 3D API, then only a fraction of them could be supported by the developer. This is the case with other APIs such as Metal (used on

Macs) and DirectX. Vulkan aims to be cross-platform. This allows the games to be easily "ported" to another platform. In other words: having this is good for games and other high-end graphical applications (e.g. medical imaging) coming to the Pi.

OpenGL ES itself is a subset of OpenGL, with a reduced feature set, specifically targeted at mobile and embedded platforms. Again, this being a cross platform standard, it brings more applications to the Pi. One significant example: WebGL. WebGL is based on OpenGL ES 2. Browsers implementing WebGL allow you to play games in the browser - without having to download and install them. And when I say games, I also mean business applications - for example the AutoCAD web app. ("*Work in AutoCAD anytime, anywhere*"). This being possible on the Pi (probably needs proper software support!) would be **huge**. It's the future of software!

An example WebGL application:

- https://www.artstation.com/artwork/B1LKk

Is hardware encoding supported for HEVC / H.265?

No, it is not supported.

Do I need a heatsink to play back 4K resolutions?

A hardware block does the decoding on the Pi 4. The Pi 4 will run a bit warmer when decoding 4K, but it should be capable of doing that without requiring a heat sink. Also note that enclosures hinder the heat dissipation through reducing ventilation and air access to the Pi 4 - depending on the design of the enclosure more or less.

Is it possible to decode two videos at the same time?

Yes, it should be - similarly to the previous Pi's. As the hardware block is shared between the two streams, you should not expect to be able to decode two 4K HEVC streams at the same time.

Will Netflix work?

libwidevine should still work for Netflix.

Will 4K resolutions work for H264 decoding?

No, H264 decode is limited to 1080p.

SD CARD

The bandwidth bottleneck to the SD card has been lifted somewhat, the data rate has been doubled. This should allow faster boot-times and applications being more snappy in general.

INPUT POWER: THE USB C PORT

A big reason for the Raspberry Pi Foundation persisting for quite a while with the well-known microUSB connector were the price differences between USB-C connectors and microUSB connectors.

On the Pi 4 power is supplied through a USB-C connector, and it is recommended to use a 3 A power supply. This requires either a new power supply - the Raspberry Pi Foundation has a nice official one - or an adapter from microUSB to USB C. The previous generation official Raspberry Pi Foundation supplies should work fine with this adapter, as long as downstream USB peripherals will consume less than 500 mA in total.

The official Raspberry Pi 4 USB C power supply. Available in black and white, EU / US / UK / AUS plugs.

Another interesting fact: the built-in USB 2.0 port is still present on the Pi 4's BCM2711 SoC, and it is wired up to this USB C socket. This might yield some interesting solutions in the future. For instance, you could use this port in a Linux "USB gadget" mode, and connect to your Pi via virtual Ethernet (g_ether gadget) from your

desktop computer, while simultaneously powering it. The Pi Zero W / Pi Zero already support this. On the 3B+ and other Pi's it was not possible, due to a USB hub being attached on the USB 2.0 port.

Ref: https://www.kernel.org/doc/html/v4.16/driver-api/usb/gadget.html

Another (future) possibility is to boot the Pi using the USB C port.

The Pi 4 uses a power management chip (PMIC): MXL7704-P4. The PMIC allows a fine-grained, precise voltage control, and communicates with the SoC using I2C for dynamic voltage scaling, status monitoring and sequencing control. A PMIC was first introduced on the Pi3B+, another advantage besides better control being the significantly lower part count required for the step-down converters. The Pi3B+ uses the MXL7704-R3, according to the Raspberry Pi Forum.

Ref:

- https://www.maxlinear.com/ds/mxl7704.pdf
- https://www.raspberrypi.org/forums/viewtopic.php?t=238779

Q & A

Why is a power management chip required?

The Pi internally needs additional voltages to the 5 V provided by the USB C power supply. It requires 3,3V (GPIO) / 2,8 V (TV DAC supply) / 1,8 V (SoC) and possibly also 1.1 V (LPDDR4 RAM, as per the spec.).

For all of these power supplies traditionally individual power rails with their own step downs were present. The PMIC integrates the necessary step downs and supporting components into one package, and adds power management logic functions.

What are the ratings for the USB C port?

The port is rated only for 5 V 3 A. No additional USB-C power modes are supported.

Will I be able to draw more power from the USB ports than on the Pi 3B+?

No, the current limit is set to 1,2 A - like on the Pi 3B+. Use a powered USB hub if your application requires a higher current.

PRODUCTION LIFETIME & INDUSTRIAL DESIGN-IN'S

Due to modular compliance certification of the dual-band wireless LAN you can design-in the board into your product with significantly reduced wireless LAN compliance testing. Thus speeding up your time-to market and lowering your costs.

The Raspberry Pi 4 Model B will stay in production for nearly seven years **until at least January 2026** - probably longer, if there is still industrial demand.

The Raspberry Pi Foundation has discontinued only five products so far:

- the original Raspberry Pi Model 1 B (superseded by 1B+)
- the Pi 1A (superseded by Pi 1A+)
- the Broadcom-based official Raspberry Pi WiFI dongle
- original Pi camera (sensor was end of life)
- original Pi camera NOIR (sensor was end of life)

Now is a good time to design-in the Raspberry Pi 4 into your next Raspberry Pi project! And we (pi3g e.K.) are here to support you with that - hardware design, software development, consulting.

APPLICATIONS & FUTURE DEVELOPMENTS

Raspberry Pi 4 opens the door to some interesting applications, which I will expand upon briefly:

WINDOWS 10 ON ARM

Microsoft has a version of Windows which runs on ARM powered PCs. This time Microsoft provides backwards-compatibility with x86 applications by emulating them.

For many use cases this is good enough. Microsoft Word and Excel will run. AAA Games and other more challenging software will not. Drivers have to be rewritten for ARM / for the specific chip.

Microsoft might release a version tailored to the Raspberry Pi, as with the Pi 4 it has become powerful enough to run a desktop operating system at speeds which modern users expect. With the x86 support as outlined above, this version would be something many users might find interesting.

DISKLESS WORKSTATIONS

Using the PoE capabilities of the Pi this could even be as simple as an Ethernet and one or two HDMI connections. The Pi boots using the network, and thanks to the fast Gbit Ethernet port, and its big RAM can load applications and data on the fly. Users log in using their credentials and can move between workspaces. Imagine a call center, based on Raspberry Pi technology for the (dual) screens!

Boot over network adds some advantages:

- no more data corruption on the SD card
- if the Pi gets stolen, the data stays with you
- the system state could be reset to known states at boot, always serving up the same system
- the backend could implement data backup, etc.

A couple of advantages over x86 technology: less power consumption, no noise, lower cost.

Compared to thin client technology these diskless workstations also lower the requirements on the backend - the servers just have to serve files, which will be cached in the Pi's RAM. The Pi will be doing the computations.

BROWSER-BASED WORKSTATIONS

Closely related to the idea above; Even more radical: Applications are not served locally anymore, but served from the Web. With the true Gigabit Ethernet port, the Pi can really benefit of broadband speeds. Using the faster CPU you will be able to use web applications, e.g. at Point of Sales systems, in the warehouse, to display Digital Signage. Without doing any special set up - just type a URL and use the application. The world is moving this way.

These applications can be rich in multimedia, as the Pi 4 supports H.265 and accelerated graphics rendering.

You will be able to work on a dual-screen system, which boosts productivity significantly.

RASPBERRY PI CLUSTERS

The next supercomputer will not be built out of Raspberry Pi's - in this area, x86 cores are still dominating the market, and smaller manufacturing nodes need to be used.

BUT - and a very significant but, the Pi 4 will be incredibly useful for small-scale clusters. For teaching purposes (Docker etc), but also as a "poor man's cluster". For instance it could be used in remote villages where you want to have a fault-tolerant system.

COMPUTER VISION

The new more powerful GPU will allow more computer vision applications & interactions - in manufacturing, for instance.

THIN CLIENTS

Specifically with the two supported displays, and the Gbit Ethernet, the Pi is now more than powerful enough to work as a thin client. A thin client does not run an operating system, it merely forwards a remote computer's output on to the local display, and the user's input and USB devices to the remote computer.

This might enable organizations to deploy some servers, for instance using Windows Server, and many cheap Pi 4 for the users. If a Pi gets broken, you exchange it for a new one. The Pi's GPIO port will enable some unique twists to the thin client idea, for instance for authentication using NFC, etc.

NAS / OWNCLOUD

Many users actually use the Pi 3 as a NAS. With the Pi 4 a NAS or Owncloud kind of solution goes mainstream. The Pi 4 should be able to saturate the Gigabit Ethernet, serving your files from an attached SSD or hard drive.

With its 4 GB of RAM and its rich Raspbian ecosystem it should enable a bigger choice of features and customizations than off-the-shelf NAS systems, which frequently are quite limited by their hardware.

MINI SERVER

Why have just a NAS when the Pi can do so much more without breaking a sweat?

- DNS server
- database server
- local Wikipedia mirror
- Wordpress server

In fact, using a Gbit Ethernet adapter on the USB 3 port, you can even build a kind of router with it to filter Internet traffic.

I can even imagine a Pi 4 (or a small Pi 4 cluster) being the central part of a (physical shipping) container, monitoring the container and providing data services.

AND MUCH MORE

We have just begun to explore the possibility space which the Pi 4 opens up for us. The new features allow for a broad range of ideas, projects and industrial applications to take form.

We are here to support you with your (industrial / business) projects - get in touch with us for a discussion of what the Pi 4 can do for your organization today!

support@pi3g.com

MARKET OVERVIEW: WHERE IS THE COMPETITION?

An important point of advice when considering a competitive product is to check it's software support - are the specified features of the SoC also supported by Linux / Android / whatever you want to run on it?

Furthermore, consider whether an active community is involved with this particular single board computer, whether there are how-tos, example projects, and whether compatible extension hardware exists.

As a rough guide to the following list, there are two big groups of SBCs:

The first group are the x86 based ones, which are still more powerful, have dual 4K outputs, support gaming and have better GPUs etc. - but also more expensive, beyond $ 100 per unit. Most often you will have to add your own RAM to these boards, making them even more expensive, but more flexible.

The other group are ARM SoC based ones. This is a big range of SBCs, including, of course, the Pi 4 itself. Other manufacturers often use Amlogic SoCs (e.g. the S922X), and Rockchip SoCs (e.g. RK3399) for their high-end models. These boards are closer to the Raspberry Pi price tag, and often add interesting additional features (e.g. PCI Express port). For some projects, they might be a better fit.

UDOO BOLT

The Udoo BOLT v8 is based on the AMD Ryzen Embedded V1605B (fabricated on a 14 nm process). This CPU operates at 2 GHz (with a boost frequency of up to 3.2 GHz).

This is a high-end x86 maker board targeted at being able to run Windows, play AAA games. With a $ 279 price tag (kickstarter) it is definitely more expensive than the Pi 4. And this does not include the RAM.

Otherwise, it matches or exceeds the Pi 4's new features, except for the WiFi:

- RAM: up to 32 GB possible (DDR4)
- 32 GB eMMC on-board
- 2 x USB 3.0, 2 x USB 3.1 Type C
- Gigabit Ethernet (1 port)
- 2 HDMI ports, version 2.0a (4k @60 fps resolution)
- **no WiFi**

Get this board if you want to play high-end demanding games on your SBC. For "traditional" maker and industry automation projects, the Pi 4 is a cheaper and more compact alternative.

Ref:

- https://www.udoo.org/udoo-bolt/
- https://www.hackerboards.com/product/263/

NVIDIA JETSON NANO DEVELOPER KIT

The NVIDIA Jetson Nano Developer Kit is priced at about $ 129. It includes a powerful GPU, four USB 3.0 ports and a DisplayPort (instead of the second microHDMI port as on the Pi).

The Jetson Nano Developer Kit might be a better choice for applications requiring a powerful embedded GPU by one of the market leaders, in other cases the Pi 4 / 4 GB is probably the better deal - at a $ 55 price point and with the large community behind it.

ROCKCHIP RK3399 BASED DEVICES

The Rockchip RK3399 has some interesting features, which match the feature set of the Pi 4 somewhat:

- 28 nm HKMG (high K metal gate) process
- dual-core Cortex A72 @ 1,8 GHz + quad-core Cortex A53 @ 1,4 GHz = 29536 DMIPS together
    - in big.LITTLE architecture - the A72 are optimized for performance, and the A53 are optimized for low power
    - the Pi 4 has four Cortex A72 cores @ 1,5 GHz = 28320 DMIPS
- Mali-T864 GPU, OpenGL ES1.1/2.0/3.0/3.1, OpenCL, DX11
- 4K 60 Hz output HDMI 2.0
- DisplayPort 1.2 (4 lanes, up to 4K 60 Hz)
    - on the Pi 4 we have dual HDMI
    - only one interface at a time is apparently capable of 4K output on this SoC - the other is driven with 2K or lower resolution (VOP_LIT).
- DRAM controller supports LPDDR4
- PCIe 2.1 with 4 full-duplex lanes
    - on the Pi 4 we have PCIe 2.0 with 1 full-duplex lane -
    - PCIe 2.1 and PCIe 2.0 have the same speed, PCIe 2.1 supports some PCIe 3.0 features.
- 2 type-C USB PHY embedded, supports DisplayPort 1.2 Alt mode on USB Type-C
- 2 USB OTG 3.0 interfaces embedded
- 2 USB 2.0 Host interfaces embedded
- embedded GMAC (10/100/1000M Ethernet Controller)

Seeing as the chip has two independent USB 3.0 controllers, and two USB 2.0 controllers in addition to the PCI Express interface, the theoretically achievable bandwidth will be higher than on the Pi 4. It will matter very much how the chip is implemented internally, and whether the CPUs would be able to saturate the bandwidth. Some manufacturers of Rockchip RK3399 based devices implement a PCI Express port, or an M.2 interface. The Pi 4 lacks those.

Here's a list of Rockchip RK3399 based devices:

- Pine64 ROCK Pro 64 (includes 1 x USB C host with video out, and PCI Express Port, no WiFi) ~ $ 80
    - https://wiki.pine64.org/index.php/ROCKPro64_Main_Page
    - https://store.pine64.org/?product=rockpro64-4gb-single-board-computer
- Firefly-RK3399 (includes eMMC) - $ 259
    - http://en.t-firefly.com/product/rk3399.html
- Rock960 (96rocks) Model B (4 GB LPDDR3, includes 32 GB eMMC, no Ethernet port, M.2) - $ 139
    - ROCK960 (96rocks) Model C (4 GB LPDDR4, no eMMC, no Ethernet, M.2) - $ 99
    - https://www.96boards.org/product/rock960/

- OrangePi RK3399 (only 1 x USB 3.0 Type C, 4 x USB 2.0 Host, Dual Band Wireless) - $ 89
  - http://www.orangepi.org/Orange%20Pi%20RK3399/

A more specific example of a (barely) sub-100 € SBC is the Radxa ROCK Pi 4 Model B, with 4 GB RAM and WiFi.

This SBC is based on the **Rockchip RK3399**, and has the following features:

- Mali T860 MP4 GPU
  - supports OpenGL ES 1.1/2.0/3.0, OpenCL 1.2, DirectX 11.1
- H.265/HEVC 4K decode
- 4 GB LPDDR4 RAM @ 3200 Mb/s
- eMMC optional 8 GB - 128 GB
- microSD slot
- M.2 SSD connector
- **one** HDMI port
- 802.11 ac wifi, Bluetooth 5.0
- Gigabit Ethernet with PoE support
- two USB 3.0 ports (one capable of OTG), apparently independent

Refer to these pages for more information about the Rockchip SoC:

- http://opensource.rock-chips.com/wiki_RK3399
- http://opensource.rock-chips.com/images/6/60/Rockchip_RK3399_Datasheet_V1.6-20170301.pdf

These boards all are significantly more expensive than the $ 55 Raspberry Pi 4 / 4 GB. (Keep in mind that some of the models have built-in eMMC, though). Again, because of the bigger community and excellent software support behind it, the Pi 4 is the safer and cheaper choice for your next SBC project.

BANANA PI

The most recent Banana Pi product is the BPI-M64 / M64-R18. It has the following features:

- 1,2 GHz quad-core ARM Cortex A53, SoC by Allwinner
- 2 GB DDR3 SDRAM
- 8 GB eMMC
- 1 GBit/s Ethernet
- WiFi & Bluetooth
- 1 HDMI port

It's available at around $ 56 - close to the Pi 4/4 GB price, but very far from its capabilities.

ODROID

Odroid-H2 is a $ 111, Intel Celeron J4105 (14 nm) based SBC. The RAM is not included in the price, it has two SO-DIMM slots for up to 32 GB of RAM.

Other features:

- 4 x PCIe 2.0 (M.2 NVMe slot)
- 2 x GBit Ethernet
- 2 x SATA 3.0
- HDMI 2.0 & Displayport 1.2 4K/60 Hz Video outputs
- 2 x USB 3.0
- 2 x USB 2.0

This board is more than twice the price of the Pi 4 Model B / 4 GB, but has some interesting features (like Dual Ethernet, and Dual 4K/60 Hz video outputs, on-board SATA, independent USB 3.0 ports for more total bandwidth) which might be interesting for high-end SBC applications. Or you might just use two Pis ...

Odroid-N2 is a **$ 79** based SBC which came out in 2019. Some interesting features:

- 12 nm based **Amlogic S922X SoC**, quad-core Cortex-A73 @ 1,8 GHz + dual-core Cortex-A53 @ 1,9 Ghz
    - the 12 nm technology allows 1,8 GHz operation without thermal throttling for the Cortex-A73 cores
- it is mounted on a big heatsink, allowing for quiet operation without fans
- 4 GB DDR4 RAM
- Gbit Ethernet
- 1 x HDMI 2.0
- 4 x USB 3.0 (via a GL3523 USB 3.0 hub)
- 1 x micro USB 2.0 OTG
- eMMC module socket
- RTC backup battery connector
- **no Wifi on-board**

The Cortex-A73 serves as the successor of the Cortex-A72, with 30 % increased performance or 30 % lower power consumption. The 12nm technology is a very interesting (albeit, as explained before more expensive) manufacturing option to make the SoC more power efficient, thus requiring less cooling - this is the way more computing power can be delivered to the user.

Seeing as this board clocks the CPU cores higher, has a smaller manufacturing node, and two additional Cortex-A53 cores, it will outperform the Pi 4. Aside from the slightly bigger form factor, and the lack of WiFi, this is a very interesting alternative to the Raspberry Pi 4 for people requiring even more performance. Please note that this board only has one HDMI port. If you require two HDMI ports, go for the Pi or look at the Odroid-H2.

Ref:

- https://forum.odroid.com/viewtopic.php?t=33781
- https://www.hardkernel.com/shop/odroid-n2-with-4gbyte-ram/
- https://wiki.odroid.com/odroid-n2/hardware

CLOSING CONSIDERATIONS

Providing a significant technology leap and keeping the 35 $ price tag is a massive engineering achievement. This becomes even more astonishing if you take inflation into account: 35 $ in 2012 are the equivalent of 39,04 $ in 2019 when adjusted for inflation and buying power.

So the Pi is not only becoming better, but actually cheaper each passing year!

I hope that you will enjoy the new Raspberry Pi generation, the one that opens the door to true desktop-class computing.

And remember - we (pi3g) are here for you on your journey to discover the Pi 4, and act as guides to Aladdin's cave of wonders and the riches it might hold.