

DPI

*De* Programmatica *Ipsium*

DE PROGRAMMATICA IPSUM

---

# **Issue 014:**

# **Minimalism**

November 4th, 2019

# Table of Contents

Issue 014: Minimalism .....	5
Do I Need A Supercomputer? .....	9
Complex Vs. Complicated .....	15



# Issue 014: Minimalism



November 4th, 2019

Welcome to the fourteenth issue of *De Programmatica Ipsum*, dedicated to the subject of *Minimalism*. In this edition:

- Graham argues that we do not need more computing power<sup>1</sup> than that of a 25 dollar computer nowadays.
- Adrian explains why “complex” is better than “complicated.”<sup>2</sup>

Enjoy this issue! Please subscribe to our free newsletter<sup>3</sup> to stay updated about new releases.

ISSUE 014: MINIMALISM

Cover photo by Igor Son<sup>4</sup> on Unsplash<sup>5</sup>.

REFERENCES

<sup>1</sup> <https://deprogrammaticaipsum.com/do-i-need-a-supercomputer/>

<sup>2</sup> <https://deprogrammaticaipsum.com/complex-vs-complicated/>

<sup>3</sup> <https://deprogrammaticaipsum.com/newsletter/>

<sup>4</sup> [https://unsplash.com/@igorson?utm\\_source=unsplash&utm\\_medium=referral&utm\\_content=creditCopyText](https://unsplash.com/@igorson?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText)

<sup>5</sup> [https://unsplash.com/s/photos/minimalism?utm\\_source=unsplash&utm\\_medium=referral&utm\\_content=creditCopyText](https://unsplash.com/s/photos/minimalism?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText)



# Do I Need A Supercomputer?



By Graham Lee

When Apple introduced their new line of Power Mac G4 computers in 1999, they famously made much of their classification<sup>1</sup> as a supercomputer<sup>2</sup> under the export controls legislation in the United States at the time. Processors (whether CPUs, GPUs, or application-specific circuitry like TPUs or the T2 security processor in modern Macs) have become both faster and more parallelized since then. Many of

the chips in a modern smartphone are much more powerful than the Motorola PowerPC 7400 that Apple claimed was weapons-grade.

A feature that Power Mac G4 CPUs, modern smartphone processors, and cores in the supercomputers on the TOP500<sup>3</sup> list have in common is that they spend most of their time waiting. However, they do this waiting for different reasons.

Ken Batcher, an electrical engineer who designed supercomputers for Goodyear Aerospace and discovered parallel sort algorithms, quipped that supercomputers were tools for turning compute-bound problems into I/O-bound problems. In other words, you put so many processors into a supercomputer that the difficulty becomes getting the data between the storage and the processors. The maths is the quick part.

These days, it's fairer to say that supercomputers are tools for turning compute-bound problems into memory latency-bound problems. Not only have processors (of any strip) become faster, but multiple cores, SIMD (Single Instruction, Multiple Data) processing and other techniques have multiplied out the number of instructions that can be running at once in a single silicon package a hundredfold or more. You can easily put a few terabytes of RAM into a single node of a supercomputer cluster, then have to wait for the data you need to get all the way from there through a rat's nest of different caches to the logic unit in the chip. Any failure in the cache control logic—or correct functioning of a caching strategy that doesn't suit your code—increases the amount of waiting.

One thing supercomputers don't tend to wait for is problems to solve. The high amount of capital investment, operational expenditure, and rapid depreciation of high-performance computing assets all mean that owners are keen to extract maximum utilization out of them. Clusters frequently expose resource-management middleware such as SLURM<sup>4</sup> which maintains queues of days or even months of pending jobs.

Meanwhile, waiting to be given a task is much of what a modern personal computer (of any size) does. A phone spends most of its time in its owner's pocket, and while a dedicated baseband processor keeps the connection to the cellular network alive the main CPU and the graphics processor have little to do. Advances

in semiconductor technology over the years have improved the potential to switch all or part of these processors off most of the time, to save energy.

Desktop computers spend just as much time doing nothing. The online game type racer<sup>5</sup> tells me that on this keyboard, I can touch type at 71wpm. That corresponds to 355 keypresses per minute, approximately six per second, or approximately 320 million CPU cycles per keypress. Every few million cycles, the CPU has to take a character code from the USB HID driver and turn it into a key event to dispatch to the browser application. Every sixtieth of a second, the GPU has to go to the great effort of working out whether, and how, to draw another character on the screen. Sometimes the characters are spaces. And that's just when I'm typing. If I'm thinking, the computer is taxed by blinking the I-bar every second.

Ah, but modern operating systems do so much more than earlier ones, don't they? Not really. It's true that there are many more features in a modern operating system, but that doesn't mean that they've necessarily expanded to make full use of modern core counts or clock rates. If I open a task monitor on the computer at which I'm editing this post, I see 10% CPU use. There are four real cores (eight virtual cores), so less than one CPU is taxed by "waiting". In fact, most of that reported use is accounted for by the activity monitor itself.

So can we get by with lesser hardware? Experience says yes. I recently replaced my desktop computer with an ARM SoC<sup>6</sup>, the sort of thing you're more likely to see in a mid-range Android phone. Even on that, the browser is fully responsive on an HD screen, video plays well, and most applications work without feeling sluggish. When you "do something" like compiling or booting the computer, it's observably slower, but even I need a cup of tea every now and then. Most of the time, around 1% of the CPU's capacity is being taxed.

Granted, even this computer is much more powerful than the Power Mac G4 that first got the Clinton administration worried. The CPU is faster, it has more cores, there is more RAM, and the operating system is stored on a relatively quick eMMC. And that's the interesting part of this situation. A \$25 tiny board is way over-specified for much desktop computing use, yet still we're pushed into spending \$1000-\$3000 on even higher specs.

In theory, the “Megahertz Myth” should have died. Apple spent years telling us that we didn’t need more MHz, back when their computers had lower MHz than their competitors’. But now they tell us that each iteration of their hardware is “the fastest yet”.

In many situations, we don’t need the fastest yet, but can only buy the fastest yet. In many situations, updated software doesn’t need the fastest yet, except that it’s made by hardware companies who want us to buy the fastest yet.

We should lead the fight against unnecessary obsolescence by example. We should do it because it’s a better use of limited environmental resources, and we should do it because it includes more people in the information revolution. Replace your overpowered desktop with a \$25 SoC, or even better don’t replace it again, ever. Choose software that works on smaller computers, or older computers, or older, smaller computers. Advocate for the adoption of that software on those computers wherever possible.

Photo by vanguard on Flickr<sup>7</sup>.

## REFERENCES

- <sup>1</sup> <https://www.techjunkie.com/apples-1999-power-mac-g4-really-classified-weapon/>
- <sup>2</sup> <https://www.techjunkie.com/apples-1999-power-mac-g4-really-classified-weapon/>
- <sup>3</sup> <https://www.top500.org>
- <sup>4</sup> <https://slurm.schedmd.com/overview.html>
- <sup>5</sup> <https://play.typeracer.com>
- <sup>6</sup> <https://twitter.com/iwasleeg/status/1138811353622360064>
- <sup>7</sup> <https://flickr.com/photos/vonguard/289999533/in/photolist-rCjKX-qzfqJj-eP2Mxw-6cU17S-67nvod-a5i4z-4vi9Gg-2Rrhy-mE5W2-aF4oH5-4vndBU-f9SYrF-io1qcq-4J4dCo-8zGuqU-25L8Zab-5uTur1-cEZfVW-gz1nSz-p3mvTD-QwpEDr-5C1Ekd-9UFSap-KUDgu-482qU2-4vi9EB-bx2JeY-NR1zA-fjHcAE-23K4pUn-wRuBZ-7yfQSN-5Dd3hU-24dPWdf-3GSt2D-bYuoDS-c2on7C-6DaKF-9pFjE-ufUJ8r-aufNM-4vi9Vv-4zyr9h-bKWsjX-bKWszM-bKWvhi-QPZLe-oQLPEB-68qi1X-a55xg8>



# Complex Vs. Complicated



By Adrian Kosmaczewski

Do you know the difference between “complex” and “complicated”? Most of us certainly mix both concepts, sometimes even several times a day. Spoiler alert: they do not mean what you think they mean.

Being an etymology geek, I opened my beloved “Dictionnaire Historique de la Langue Française<sup>1</sup>,” this monumental work by Alain Rey published chez Éditions Robert, and found out the ancestors of both words. Not just because doing it in French makes everything look more intellectual, mind you.

*Complex is borrowed from the Latin complexus, meaning “made of intertwined elements”. This word is the past participle adjectived of complecti (“embrace, understand”) itself derived from plectere (“bend, intertwine”). This word has been used since the XVI century to qualify that which is made of heterogenous elements.*

Complex shares the same root (*plectere*) with the words “perplex” (in use since the 14th century in France) and the medical term “plexus” meaning “interlacing” and used since the 16th century as a medical term for “network of nerves or blood vessels”.

One of the most common uses of the word in science is the concept of “complex numbers”, created by the Italian mathematicians Gerolamo Cardano and Raffaele Bombelli (between 1545 and 1560), later rigourously defined by William Rowan Hamilton in 1828, born from the need to solve the problem of the square root of  $-1$ .

On the other hand, “Complicated” has a similar origin but a different construction:

*Complicated comes from the Latin complicare, literally meaning “to fold by rolling up”. Figuratively speaking this was taken as close to the notion of embarrassment or awkwardness. The word is composed of the word plicare which means “to fold”.*

The Spanish word “plegar” and the French “plier” are direct descendants of this Latin root.

To summarise, “Complex” and “Complicated” stem from slightly different roots: the Latin root *plectere* (“to intertwine”) in the former, and *plicare* (“to fold”) for the latter.

*Complex* conveys the idea of a network of intertwined objects, whose state and behavior are continuously modified by their peers in that network. The word *complicated* implies an intrinsic apparent “obscurity” through folding unto itself, inviting to a simple “unfolding” discovery process.

Or, put in another way: complex numbers are not complicated. They are made of real and imaginary parts. Maths can be complicated, though.

In page 12 of the “NeXT Object-Oriented Programming and the Objective-C Language<sup>2</sup>” book, available at the GNUStep project website<sup>3</sup>, the image of a pocket watch conveys an explanation to the abstract concept of “encapsulation.” Encapsulation of what? Of complication, precisely.

Watches commonly known as “Complications” (such as the Patek Philippe Calibre 89<sup>4</sup>, the Franck Muller Aeternitas Mega<sup>5</sup> and the “Référénce 57260” de Vacheron Constantin<sup>6</sup>) are complicated, yet not complex, machines. An “Object” in the sense of OOP, should a complicated, but not a complex, thing. Not exactly the experience reported by most code reviewers I have met so far. (By the way, these days the word “complication” is still associated to watches, albeit meaning something entirely different<sup>7</sup>.)

The Kubernetes-fueled *microservices* craze has everyone rewriting complicated monoliths into complex cloud-native apps. As my partner in crime Graham once famously wrote in his opus “OOP The Easy Way<sup>8</sup>,” microservices is nothing but a new way to say object-orientation. Loosely-coupled entities communicating with each other with messages, anyone?

Simple entities sending messages to each other. Like living cells using mRNA<sup>9</sup> to transfer protein pattern information from the DNA to the ribosome. Uncomplicated, minimalistic entities, used to build larger, more complex systems. Have we forgotten about that? Do we need to learn this over and over?

Software developers have a love / hate relationship with complicatedness, and an almost instant disdain for (if not outright ignorance of) complexity. Complicated systems are great to brag about on Reddit; maintainers also cry about them in private.

A “best practice” has one and only one basic job: to help engineers translate complicated into complex. Most software-related disasters are caused by a simple fact: because of deadlines, managers, tooling, or just plain ignorance, software developers have a tendency to build complicated, instead of complex, systems.

Let us write and run complex, not complicated, systems. We cannot get away from complexity; that is our job as engineers. But we can drop the complicated bit, and minimalism will follow.

Instead of running a 500 MB Docker container with Linux, Apache, PHP, and the full source code of that monolith written in 2002, let us instead run a 5 MB container built FROM scratch running a microservice written with Go<sup>10</sup>.

Instead of running Windows, macOS or Ubuntu, let us boot our PCs with KolibriOS<sup>11</sup>, Tiny Core Linux<sup>12</sup> or Damn Small Linux<sup>13</sup>. Or even with Minix<sup>14</sup>, MikeOS<sup>15</sup>, or FreeDOS<sup>16</sup>.

Instead of writing mobile apps with yet another hyped layer of abstraction, use the lower-level language recommended by the platform vendor, giving you direct access to the native APIs of the device, and a tighter control of the memory and CPU consumption of your app.

Instead of using yet another programming language several dozen abstraction layers away from the actual platform, let us go back down a bit, thanks to languages like Go<sup>17</sup>, Rust<sup>18</sup>, D<sup>19</sup>, Crystal<sup>20</sup>, or Elixir<sup>21</sup>. Thanks to projects like LLVM<sup>22</sup>, many high-level programming languages do not need need a runtime layer anymore, can be cross-compiled, and offer blazingly fast performance.

Instead of Emacs, let us try Joe's Own Editor<sup>23</sup>.

And after having launched it, let us write our own operating system<sup>24</sup> with it.

Cover photo by Dan Lohmar<sup>25</sup> on Unsplash<sup>26</sup>.

## REFERENCES

- <sup>1</sup> [https://fr.wikipedia.org/wiki/Dictionnaire\\_historique\\_de\\_la\\_langue\\_fran%C3%A7aise](https://fr.wikipedia.org/wiki/Dictionnaire_historique_de_la_langue_fran%C3%A7aise)
- <sup>2</sup> <http://www.gnustep.org/resources/documentation/ObjectivCBook.pdf>
- <sup>3</sup> <http://www.gnustep.org/>
- <sup>4</sup> [https://en.wikipedia.org/wiki/Patek\\_Philippe\\_Calibre\\_89](https://en.wikipedia.org/wiki/Patek_Philippe_Calibre_89)
- <sup>5</sup> <https://www.franckmuller.com/aeternitas/>
- <sup>6</sup> <http://reference57260.vacheron-constantin.com/fr/la-montre-la-plus-compliquee-au-monde>
- <sup>7</sup> <https://www.imore.com/how-add-third-party-complications-your-apple-watch>
- <sup>8</sup> <https://leanpub.com/ooptheeasyway/>
- <sup>9</sup> [https://en.wikipedia.org/wiki/Messenger\\_RNA](https://en.wikipedia.org/wiki/Messenger_RNA)
- <sup>10</sup> <https://golang.org/>
- <sup>11</sup> <http://kolibrios.org/>
- <sup>12</sup> <http://tinycorelinux.net/>
- <sup>13</sup> <http://www.damnsmalllinux.org/>
- <sup>14</sup> <http://www.minix3.org/>
- <sup>15</sup> <http://mikeos.sourceforge.net/>
- <sup>16</sup> <https://www.freedos.org/>
- <sup>17</sup> <https://golang.org/>
- <sup>18</sup> <https://www.rust-lang.org/>
- <sup>19</sup> <https://dlang.org/index.html>
- <sup>20</sup> <https://crystal-lang.org/>
- <sup>21</sup> <https://elixir-lang.org/>
- <sup>22</sup> <http://llvm.org/>
- <sup>23</sup> [https://en.wikipedia.org/wiki/Joe's\\_Own\\_Editor](https://en.wikipedia.org/wiki/Joe's_Own_Editor)
- <sup>24</sup> <http://mikeos.sourceforge.net/write-your-own-os.html>
- <sup>25</sup> [https://unsplash.com/@dlohmar?utm\\_source=unsplash&utm\\_medium=referral&utm\\_content=creditCopyText](https://unsplash.com/@dlohmar?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText)
- <sup>26</sup> [https://unsplash.com/s/photos/complex?utm\\_source=unsplash&utm\\_medium=referral&utm\\_content=creditCopyText](https://unsplash.com/s/photos/complex?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText)