

What every research software engineer should know about floating point and one bit more

Floating point numbers are the backbone and deepest essence of computation. Floating point numbers is how our computer represents and manipulates “real numbers” (wanna know why the quotation mark? come to the talk). Unfortunately, floating point arithmetic is considered too applied by mathematicians and too theoretical by engineers, hence it often neglected in computer science or scientific computing curricula, and considered an obscure and dry field.

That’s not true! Floating point arithmetic is a fascinating realm, where all the maths you have learnt is to be questioned again. Floating point arithmetic is not “wrong or approximate math” but its own maths, full of properties and theorems.

This talk is meant to be a dwarf on giants’ shoulders, a revisited and colourful Gen-Z edition of the legendary gem “What every computer scientist should know about floating point arithmetic”. In this talk, we will open the pandora box and demistify floating point numbers, discussing what they actually are, how they work and why they work. We will also give a few extra bits to get more accuracy when you need it.

After the talk, the audience will know everything they need to know to work with floating-point numbers in everyday tasks. Moreover, after the talk, you will never be able to type 0.1 in your computer without getting chills through your spine, knowing you are scratching the surface of a brave new world.

Primary author: FERRANTI, Luca

Presenter: FERRANTI, Luca

Track Classification: Track 1