Acceptance speech
21 de septiembre de 2021

# **John L. Hennessy**, awardee in the Information and Communication Technologies category (13th edition)

We are very grateful to the BBVA Foundation for this prestigious award. As my colleague, Dave Patterson, explained, the large performance advantages offered by the RISC approach were surprising. Although we had qualitative explanations for these performance gains, we did not have a solid quantitative explanation of what was happening. Thus, although many of our academic colleagues believed our results, industry was much more hesitant, which was one of the key motivations behind my efforts to commercialize our research by cofounding MIPS Computer Systems.

Unfortunately, the lack of a solid quantitative explanation also made it hard for the young company to persuade customers of the advantages of RISC and their durability. As I struggled with this dilemma I decided to try to break down the time needed to execute a program into the product of the number of instructions required, the number of clock cycles per instruction (CPI), and the length of the clock cycle. While the last of these is largely a function of the technology, the number of instructions depends on the instruction set architecture (ISA) while the clock cycles per instruction depends on the effectiveness in implementing the instructions, which of course also depends on the ISA. The key insight is that the RISC approach sacrificed a small increase in instruction count (typically 1.2-1.5 times) for a large reduction in CPI (typically 3-6 times). A landmark paper analyzing the performance of the DEC VAX 11/780, then the number one selling minicomputer in the world, showed that it had a CPI of ten—more than six times that of the early commercial RISC microprocessors. Overall, the RISC microprocessors offered a performance advantage of 3-5 times, while also requiring less hardware. Numerous papers since then have confirmed these advantages.

As I prepared to return to Stanford at the end of the leave at MIPS Computers, I was puzzled: why had we not seen these insights earlier and how had the commercial architectures gone so far astray? I realized that both the way we taught computer architecture and the way we evaluated instruction sets was flawed. We taught computer architecture as a survey course focused on comparing abstract analyses of instruction sets, and we often evaluated instructions by abstract metrics: how many different instructions, how many addressing modes, and sometimes code size. What about implementation and performance? What mattered was how fast programs would run—not these abstract metrics.

I talked to Dave Patterson and found that he was equally dismayed about the state of the field and how we were teaching computer architecture: maybe we should write a book to try to change how the field was taught and how researchers and engineers thought about building computers. That was the beginning of our long collaboration.

We intended our book to reshape the way the field was taught and structured: to bring a scientific approach to the field, just as Don Knuth's famous textbooks (*The Art of Computer Programming*) had for programming and algorithms some 20 years earlier. Thus, the name of our text: *Computer Architecture: A Quantitative Approach*. The CPU performance equation became the foundation used to analyze both instruction sets and implementation approaches. We followed the quantitative approach throughout the book. For example, we introduced Average Memory Access Times (rather than just miss rate) as the key metric for analyzing a memory system.

Starting with a rough set of lecture notes, which we used in our own courses in the Fall of 1988, and taking advantage of our sabbaticals in the Spring of 1989 , we completed a Beta version that was used in the Fall of 1989. Real-time feedback from more than a dozen professors who taught with the Beta version and over 100 other researchers and engineers enabled us to finish the draft by December 1989. The book was well received and widely adopted selling almost 10,000 copies in the first year (surpassing our publishers lifetime estimates in its first year!). Surprisingly, as many copies were sold to practicing engineers as to students.

In his preface, our renowned colleague Gordon Bell said: "The authors have gone beyond the contributions of Thomas to Calculus and Samuelson to Economics. They have provided the definitive text and reference for computer architecture and design." Happily for us, he was prescient. A number of translations were started soon after and the book is now available in more than a dozen languages and is used around the world. Despite six editions and more than 30 years, we still have a great working relationship, and while those fundamental performance equations are still there, 90% of the material in the sixth edition did not appear in the first edition! We went on to write an introductory undergraduate version, called *Computer Organization and Design: The Hardware/Software Interface*. Although we had hoped that someone else would write a great undergraduate text in the same spirit as our first book, no one did. Happily, the undergraduate text enabled us to reach an even larger audience. It is in its fifth edition and together, more than 200,000 students around the world have learned from our books. Helping students around the world learn about our field through our books has been incredibly rewarding. Many, many times (most recently today), I have had a student walk up to me and say: "Are you *the* John Hennessy?" When I reply, "yes," they exclaim: you helped me learn about computer architecture, or pass a tough course, or get excited about the field.

21 de septiembre de 2021

This brings me to why the BBVA Foundation Frontiers of Knowledge Award is so meaningful to Dave and I. First, it celebrates the early, risky approach we took in our research. The RISC ideas were both counterintuitive and controversial. It is unclear even today, whether the ideas would have gotten traction, if only one team had discovered them! This award celebrates not only those insights and the importance of research that rethinks problems when the parameters change, but also the willingness to persevere when some people think you are crazy. Equally importantly, the BBVA Foundation Award recognizes the fruits of our 30 year partnership to change the way students are taught and engineers design computers. Today, it's hard to find someone working in our field that has not had the experience of using Hennessy & Patterson or Patterson and Hennessy, and this award celebrates in a special way, that global impact we have had. My sincere thanks to the BBVA Foundation and the award selection committee, to my good friend and partner Dave Patterson, to our colleagues at Stanford and Berkeley, to my family, and to the students around the world that have graciously journeyed into the field of computer architecture with us!