

John Tsitsiklis

**INFORMS John von Neumann Theory Prize, Acceptance speech
Phoenix, Arizona, November 2018**

I am of course delighted to share this great honor, especially after looking at the list of famous past recipients. I will take this opportunity to go back into some of the history, and see how we got here.

I started my graduate work in LIDS, the Laboratory for Information and Decision Systems, at MIT, in 1980, just when Dimitri also moved there as a faculty. Among many other interests, Dimitri was already looking at distributed algorithms for problems motivated from the emerging field of data networks. On my end, I was lucky to have Michael Athans as my thesis supervisor, an inspirational personality who was pushing the field of control, and optimization in new directions, involving large scale and decentralized systems.

And so Dimitri and I started interacting, and we ended up proposing and studying versions of deterministic and stochastic gradient descent that run over a loosely synchronized, asynchronous network, coordinated through a consensus algorithm. The storyline at the time involved loosely coupled workstations, an idea that subsequently lost steam. Interestingly, the subject has now resurfaced, especially in the context of large-scale optimization for machine learning. Of course at the time no one was predicting this turn of events, but this experience reinforces the following thought. Instead of trying to tug along today's bandwagon, one can just focus on clean and generic problems, on settings that feel and taste right. And then there is always a chance that the bandwagon will some day travel your way.

Of course, our collaboration on parallel and distributed algorithms went much further. It was a very exciting and productive time – academic life was much simpler in those days. Our styles were somewhat different, but not too different, and so we complemented each other nicely.

The second major topic on which we worked together, in the mid-nineties, was approximate and simulation-based dynamic programming, better known these days as "Reinforcement learning". Here, we were blessed to work with a cohort of incredible students. On my end, there was Peter Marbach, Vijay Konda, and most important Ben Van Roy, whose insightful results and deep thinking actually drove much of the research agenda.

The bottom line is that coincidences, who you get to work with, and when, make a huge difference. Crossing paths with Dimitri was possibly the happiest coincidence in my academic life, and I am truly delighted to be sharing this award with him.