Title

The Fast and the Fabulous. Harnessing GPU Power for High-performance Life Insurance Computations

Speaker/Company

Karol Maciejewski / Mehdi Echchelh, Milliman

Abstract

Data science most often relies on heavy computations on data, either due to complexity of calculations or volume of the data, or both. In this technical presentation we want to discuss our experiences working on design and implementation of high-performance large-scale computations required in the context of life insurance ALM projections using Python, GPU and clusters. Commonly used Python data science and machine learning libraries provide simple gateways to the GPU computations that don't require GPU knowledge to use. However, if one wants to go beyond that and develop fully custom models, at least some understanding of the GPU architecture and packages allowing its utilization is necessary to truly take advantage of this technology and avoid common pitfalls. We will present the difference in the computation model of the CPU and GPU, types of problems that can be efficiently tackled using massive parallelization on GPU and clusters, Python libraries that can be used by modelers with different levels of experience and some hints on what to do and what not to do when implementing GPU computations. We will also show benchmarks of our implementations to show that speed-ups in order of magnitude of up to 1000x are possible with an appropriately chosen model and implementation architecture.

We believe our considerations and examples can be very useful to anyone working with life insurance data. While we focus on multidimensional liability and asset projections, these methods can be applied to any data engineering or data science problems equally well. We also use this opportunity to highlight key differences between developing sandboxed models for ad-hoc studies versus IT-approved production models.

Biography

Karol is a director in the Life team of Milliman and a member of the Belgian Institute of Actuaries (IABE). He has worked combining his actuarial expertise with IT-technical knowledge and passion for technical innovation for life, non-life and data analytics. For the last 10 years, he has been working internationally on ALM and life insurance modeling, focusing on model and process optimization. Most recently, he has been advocating the use of and leading projects using Python and GPU computing for custom high-performance insurance models.

Mehdi is a consultant in the Data Analytics team of Milliman. He has experience in the fields of insurance and data science, and is a member of the French Institute of Actuaries (IA). He has a strong interest in computer science and has been actively involved in actuarial modelling topics. In 2021, he joined the analytics team at Milliman Paris, where he focuses on topics such as

machine learning applied to insurance, data engineering, and data visualisation.