The fast reduced QMC matrix-vector product

Peter Kritzer
Johann Radon Institute for Computational and Applied Mathematics (RICAM)
Austrian Academy of Sciences
Linz, Austria
Joint work with Josef Dick, Adrian Ebert,
Lukas Herrmann, and Marcello Longo.

In various applications (e.g., mathematical finance or PDEs with random coefficients) one is interested in approximating integrals of the form

$$\int_D f(\boldsymbol{x}^\top A) \mathrm{d}\mu(\boldsymbol{x}),$$

for a domain $D \subseteq \mathbb{R}^s$, an $s \times t$ matrix $A \in \mathbb{R}^{s \times t}$, and a function $f : D \to \mathbb{R}$, by quasi-Monte Carlo integration rules of the form

$$Q_N(f) = \frac{1}{N} \sum_{k=0}^{N-1} f(\boldsymbol{x}_k^\top A).$$

We are interested in situations where the main computational cost of computing $Q_N(f)$ arises from the vector-matrix multiplication $\boldsymbol{x}_k^\top A$ for all N points, which requires $\mathcal{O}(N s\, t)$ operations.

It was shown by Dick, Kuo, Le Gia, and Schwab that, when using particular types of QMC rules, the cost to evaluate $Q_N(f)$ can be reduced to only $\mathcal{O}(t\, N \log N)$ operations. This drastic reduction in computational cost (as long as $\log N \ll s$) is achieved by a fast matrix-matrix multiplication that exploits the fact that for the chosen point sets the matrix X can be reordered to be of circulant structure. The fast multiplication is then realized by the use of the fast Fourier transformation (FFT).

In our talk, we will outline a different method which can also drastically reduce the computation cost of evaluating $Q_N(f)$. The reduction in computational complexity is achieved by using point sets which possess a certain repetitiveness in their components.