

About the Cover

The cover accompanies the article by Ian Sloan and Frances Kuo about the construction of lattice rules for multivariate integration (see page 1320). They present an algorithm, called “component-by-component construction”, which by a mix of number theory and linear algebra yields some beautiful pictures.

In constructing a d -dimensional N -point lattice rule, we must choose in succession d components of its generating vector. In finding one of these, all possible choices for this component of the generating vector are considered, from among the integers in \mathbb{Z}_N relatively prime to N . Multiplication modulo N for each of these choices corresponds to a certain permutation of the numbers in \mathbb{Z}_N .

So we can imagine making a multiplication table by considering horizontally the numbers from 0 to $(N - 1)$ and vertically the numbers relatively prime to N . The bottom images of the illustrations give a visual impression of this table with the numbers from left to right and top to bottom in natural ordering. Since each row of this table corresponds to one of the permutations of the N points, there are N different colors per row. A matrix with this special structure pops up in the component-by-component algorithm and in each step a matrix-vector product has to be carried out.

Surprisingly, by just doing row and column permutations, this matrix can be brought into a form which allows for multiplication in time $O(N \log N)$ for any N . This can be done in two steps. The first step, and also the step which gives rise to the nicest images, is grouping the numbers in \mathbb{Z}_N according to divisors common with N . This permutation on the columns is visualized in the middle images. If N is a power of a prime this image clearly displays a multi-resolution view of the general structure present in the first block. We have lower and lower resolution tilings of the first block on top of each other. For general N , similar but more complicated patterns appear.

In the second step we pull these smaller multiplication tables apart into multiple cyclic groups. If g is a generator for one of these cyclic groups with modulus p^α , then its elements are given by $g^k \bmod p^\alpha$ and the multiplication table has the elements $g^{k+\ell \bmod \phi(p^\alpha)} \bmod p^\alpha$, i.e., having constant anti-diagonals for k and ℓ in natural ordering. Taking $-\ell$ instead of just ℓ gives constant diagonals, making the table a circulant matrix. When we have to consider more than one generator per block then we can use the Chinese remainder theorem to obtain nested block circulant matrices. Thus by permuting the rows and the columns we can reorder the partitions so that they form nested block circulant matrices as can be seen in the top images.

It is well known that a matrix-vector product with a circulant matrix can be done in time $O(N \log N)$ using fast Fourier transforms. A similar procedure delivers a matrix-vector product with this complete matrix also in time $O(N \log N)$, being more complicated for general N than for N prime (see Nuyens and Cools in the “Dagstuhl 2004” issue of *J. Complexity*, to appear). This fast matrix-vector product is an essential ingredient for constructing large lattice rules with say $N \approx 10^8$ points.

The graphics were created using the Python package PyX .

—Dirk Nuyens

dirk.nuyens@cs.kuleuven.be




Figure 1. The three steps for $N = 2 \times 3^2 \times 5 = 90$. For general N the structure becomes more complicated. (In the top matrix the grayed out blocks denote redundancy.)




Figure 2. The three steps for $N = 2^7 = 128$. As is usual in number theory, powers of 2 (the only even prime) behave a little bit different than powers of odd primes, except for $N = 2$ and $N = 4$ which are also the usual exceptions. In this case we get an extra circulant embedding. In the top matrix the grayed out blocks denote redundancy.) Also see the cover image.