# Efficient Simulation of Long Neutron Beamlines, or How Not to Waste Time on the Computer

Isak Svensson

**Have you ever found yourself spending way too much time on the computer? If so, you are not alone–some computer programs do it, too. Physicists frequently use Monte Carlo simulation to study physical processes, but the technique eats CPU cycles like a whale eats krill. So how do we get the best possible data from the least possible computing time?**

Neutrons are a fabulous tool for studying the innards of objects, but attempting to control them is mostly an exercise in futility. Bright neutron sources take a "machine gun approach" to get neutrons where they are needed: the neutrons fly every which way, and the uninteresting ones are filtered away. It's rather crude, but it's the best we can do at the moment.

Spallation neutron sources, like the European Spallation Source in Lund, Sweden, produce their neutrons by firing high energy protons at a material rich in neutrons. This results in the release of a large number of neutrons of varying energies, up to the energy of the driving proton beam. The neutrons are moderated (slowed down) after emission, but some really high energy neutrons remain. These are troublesome because they are of no scientific use and they have an impressive ability to penetrate shielding.

Neutron beamlines are used to guide the scientifically relevant neutrons to the experimental area. The beamlines are narrow ducts, a few cm across, but very long–sometimes over 100 m. Because the neutrons emerging from the source travel in essentially random directions, exceptionally few will make it to the end of the beamline. But what if you're just as interested in the particles at the end of the guide as those near the entrance? If you studied the situation using the straightforward Monte Carlo method–sampling a large number of trials to obtain an average behaviour–you'd have to wait until the end of time to gather meaningful statistics at the end of the guide.

One possible solution is to simulate the process you *wish* you had. You make sure that the neutrons are evenly distributed across the length of the beamline, and worry about realism later! This is the idea of the so called duct source variance reduction method, and the realism of the simulation is preserved by careful alterations to the particles' statistical weights. The implementation effortlessly transports a huge number of neutrons to virtually any distance, greatly reducing the computer time spent simulating neutrons travelling in uninteresting directions. Using this method makes it possible to, for example, study the efficacy of the radiation shielding around the beamlines at great distances from the source.

While few in numbers, the really high energy neutrons can penetrate most shielding with ease and wreak havoc with both measurements and personnel. This warrants extra careful investigation of this component of the spectrum. To properly study them, pretend they are more numerous than they really are! In fact, let the neutron energies be evenly distributed across the spectrum with all energies being equally likely. Tests show that this significantly increases the number of neutrons in hard-to-reach areas of the geometry. An important application will be to study the presence of high energy neutrons around the experimental area due to *skyshine*–neutrons that at first escape the facility, only to bounce back against the air.

In all, the implementation of the duct source enables otherwise intractable studies of neutron fluxes where it is perhaps most interesting–at locations far from the source position and behind thick shielding.