Radiotherapy Dose Calculation Acceleration With CUDA

Sami Hissoiny (1), Lahcen Arhjoul (2), Jean-Philippe Gariépy(2), Benoît Ozell(1), Philippe Després(2)
(1) École polytechnique de Montréal
(2) Centre de recherche de l'hôpital universitaire de l'Université de Montréal(CRCHUM)

Introduction

Dose calculation is a central part of the treatment planning routine for patients receiving radiotherapy treatments. In a clinical environment, dose calculation needs to be accurate and fast. This accuracy is obtained by the integration of patient specific data, gathered by medical imaging modalities, and extensive beam modeling. The large data sets being manipulated result in relatively slow algorithms.
Three different dose computation methods have been accelerated by CUDA: a convolution/superposition (CS) algorithm in an external photon beam configuration, a modified heterogeneous TG-43 brachytherapy algorithm and a finite size pencil beam (fsPB) algorithm.
Our goal was to accelerate dose computation while preserving or improving the accuracy of the model.

Dose calculation: a highly **parallel** task

N Voxels → N Threads

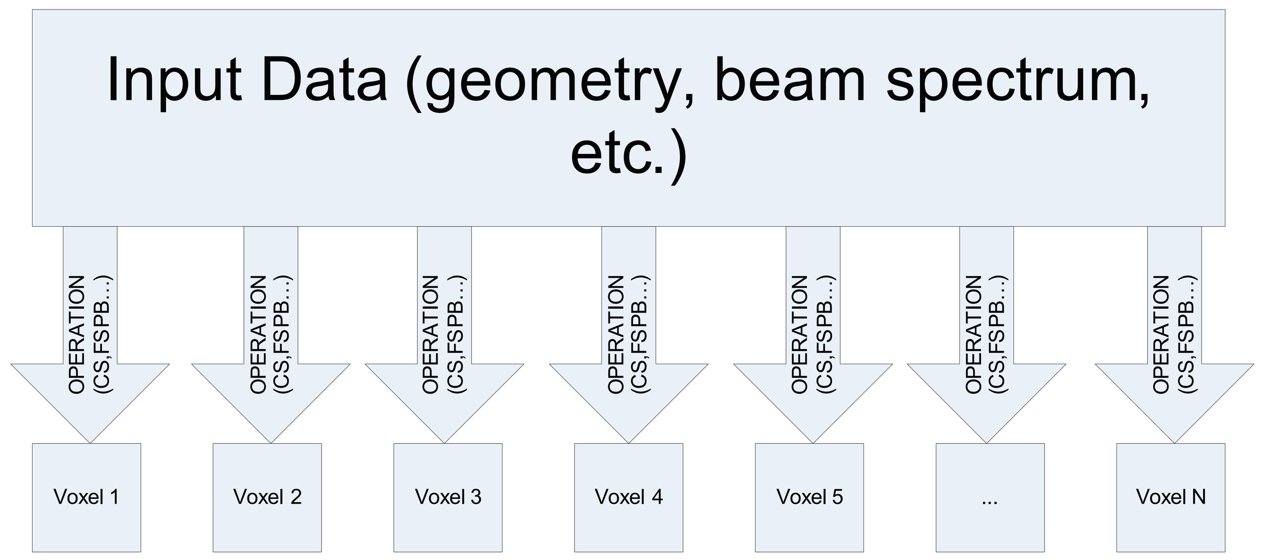

Ideal candidate for a GPU implementation

Convolution/Superposition

The CS consists in convolving the computed photon fluence coming from the linear accelerator at every voxel with an energy deposition kernel, modeling diffuse radiation. This task is computionally intensive, especially since the kernel must be scaled according to the density values contained in the computed tomography (CT) data.
The algorithm takes into consideration many of the fundamental physics phenomena involved in radiation transport but at the cost of increased computional times. Our implementation takes less than 3 seconds per beam to compute a full 3D dose, down from the several minutes required using the same algorithm on the CPU.

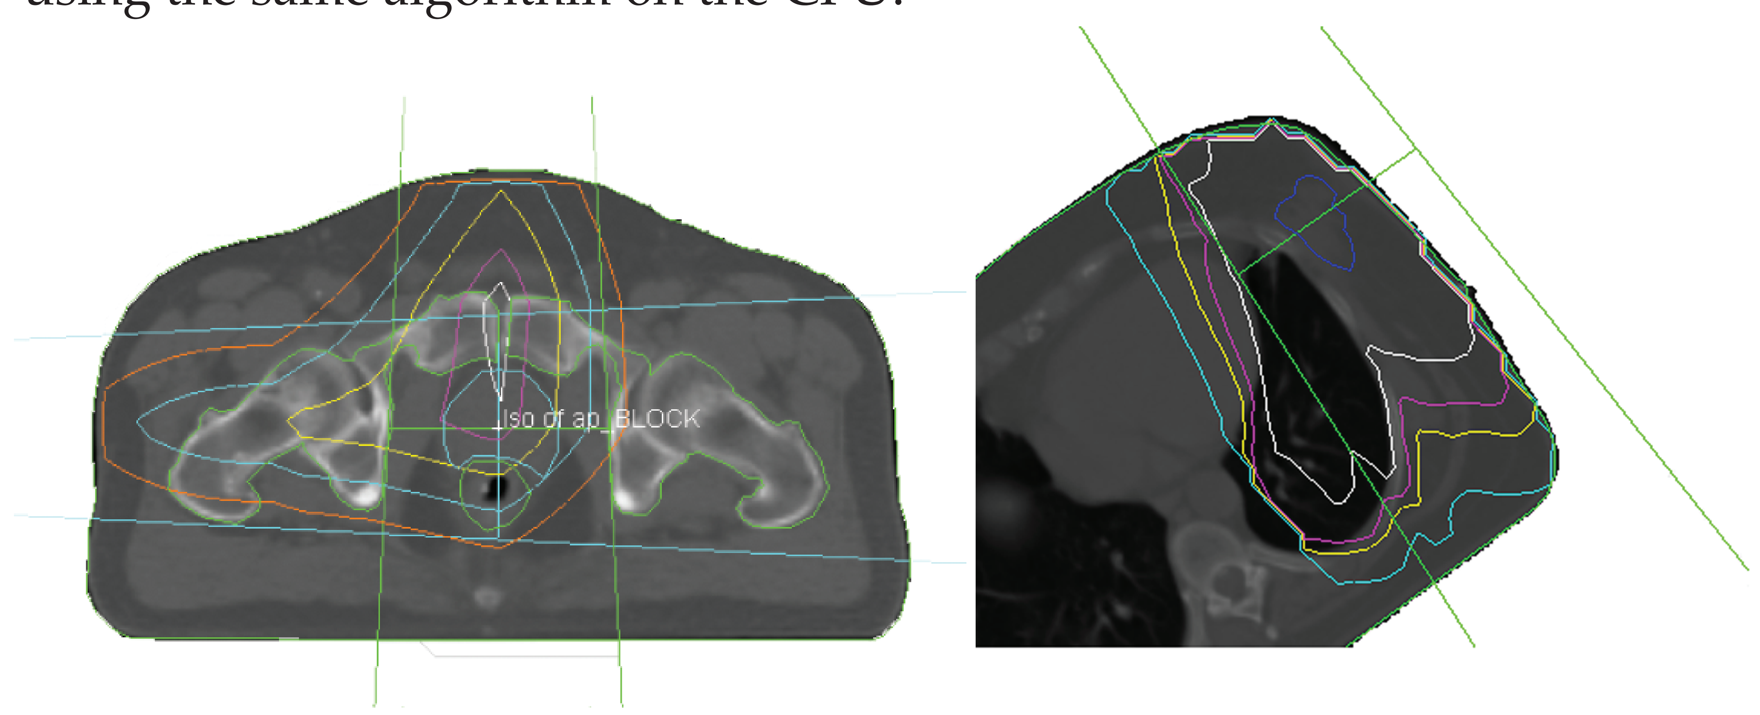

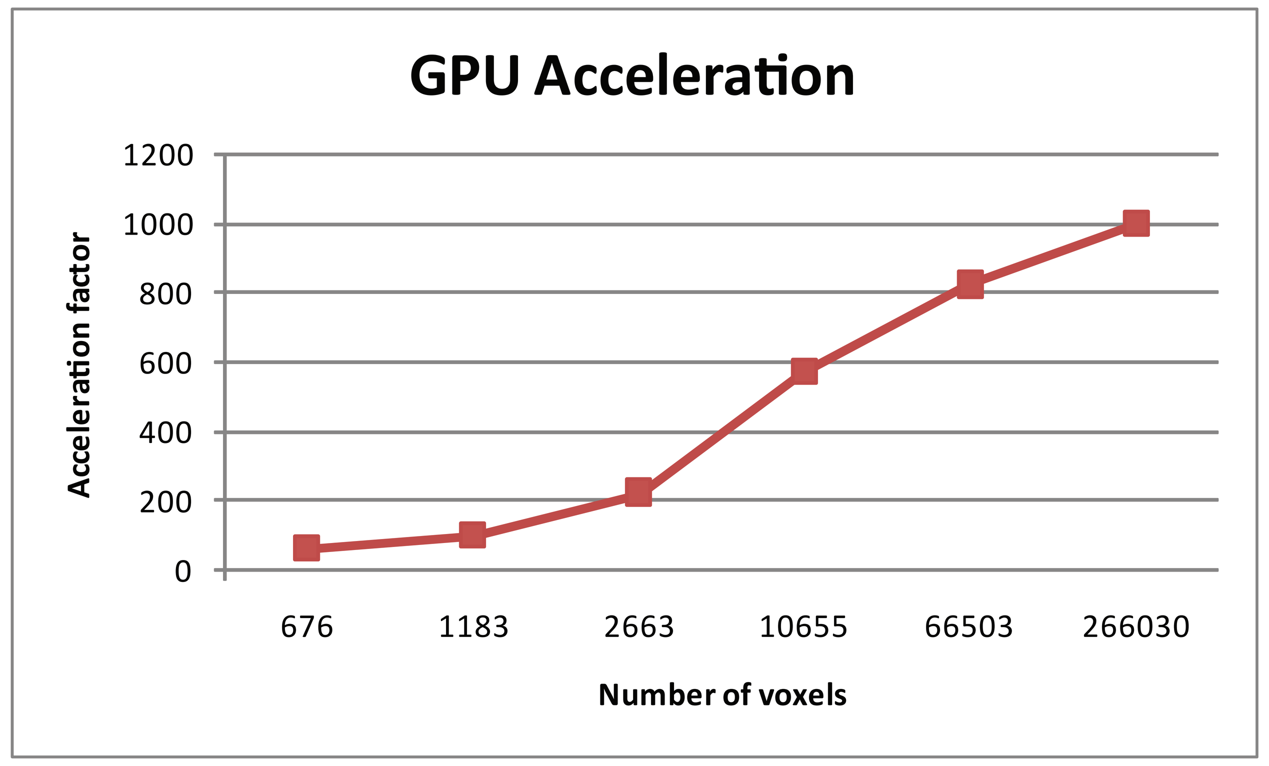

Finite Size Pencil Beam

Using this model, the dose is computed by dividing the photon beam into a number (1000s) of beamlets. The beamlets are self-consistent, *i.e.*, the sum of the subdivized beamlet is exactly equal to the original beam.
The dose at one voxel is the summation of all the beamlets' contributions. Beamlets can either be turned on or off to account for irregular fields. This allows the computation of a highly conformal dose distribution by shaping the beam to target and sparing the surrounding organs at risk.
Beamlets contributions are described as a piecewise analytical function making heavy use of trigonometric operations. The GPU implementation benifits from a high arithmetic density and low bandwidth requirements.

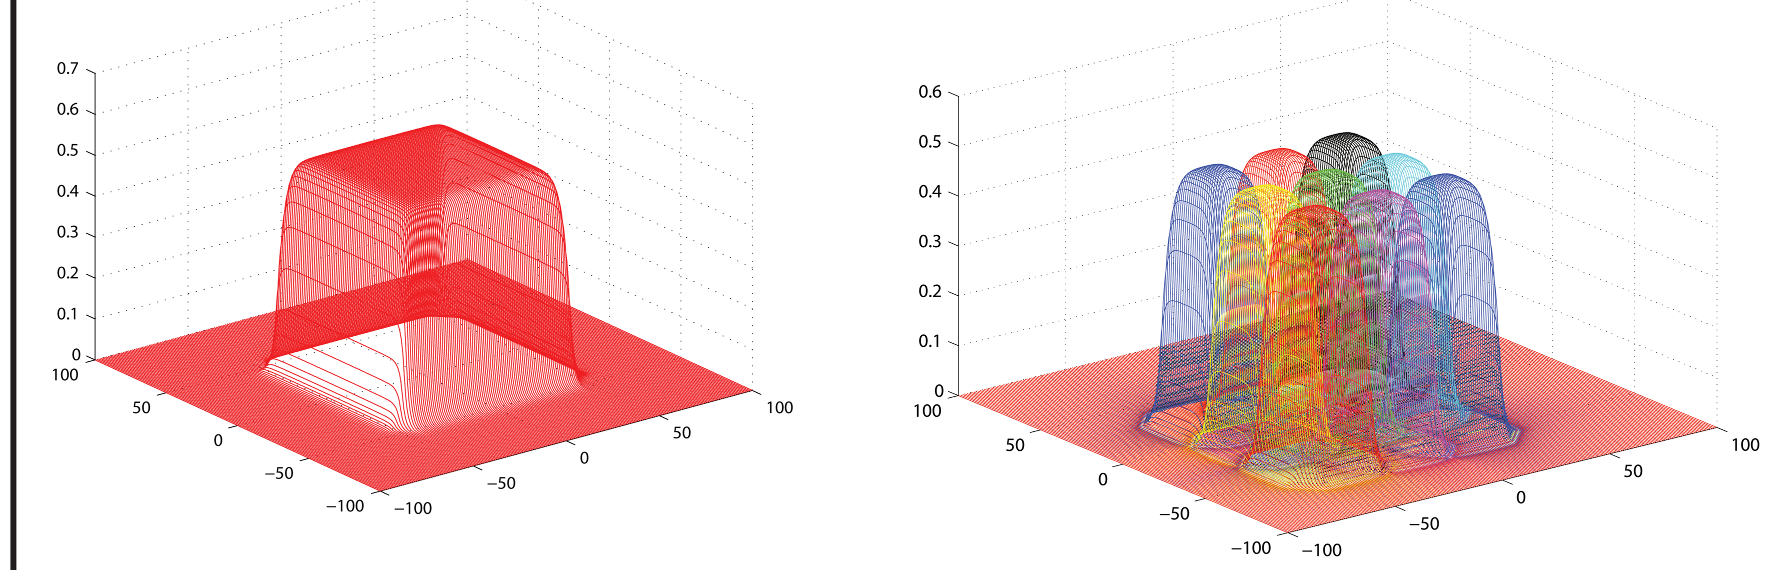

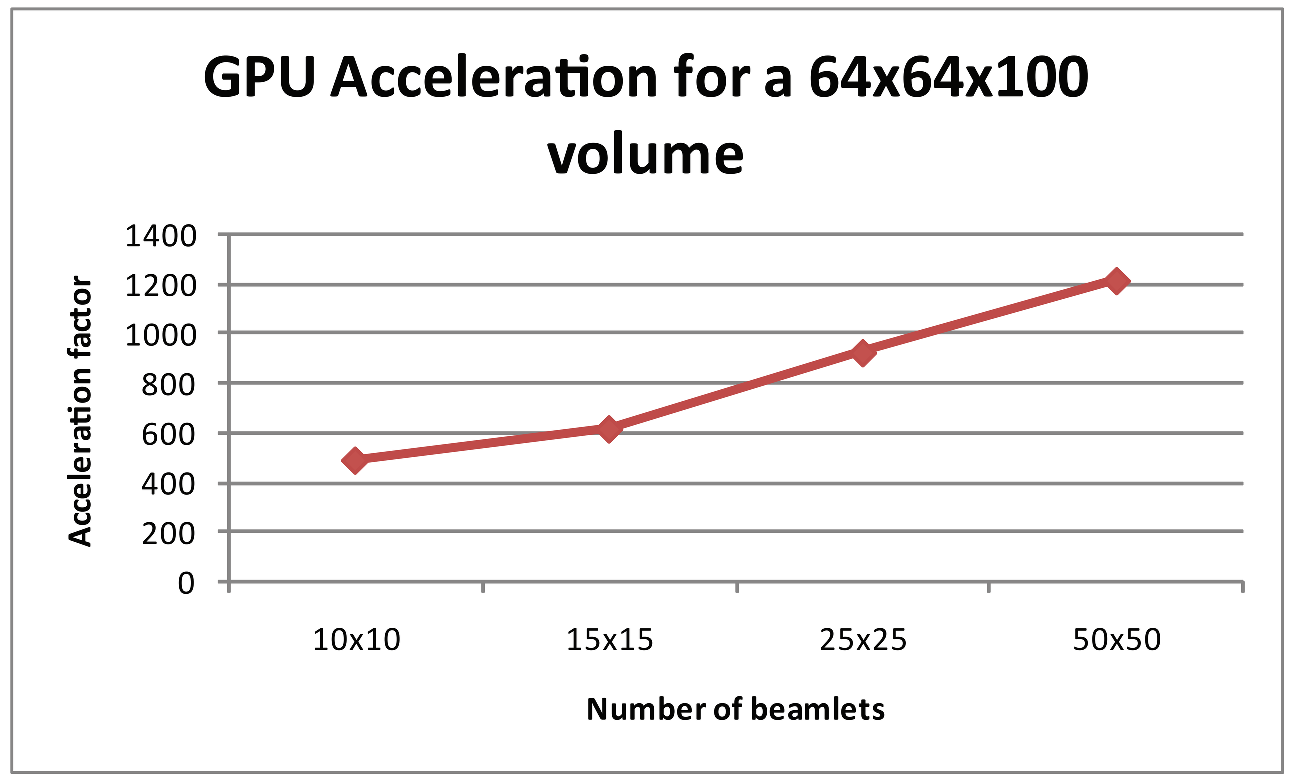

Brachytherapy

In brachytherapy, radiation sources are inserted inside a tumor site. The widely used formalism to compute the dose in this case does not take tissue heterogeneities into account.
We have implemented the formalism on the GPU and added a full raytracing engine to add tissue heterogeneity corrections. We have therefore improved to the accuracy of the model.
The same implementation running on the CPU did not produce results within a timeframe compatible with clinical workflow.

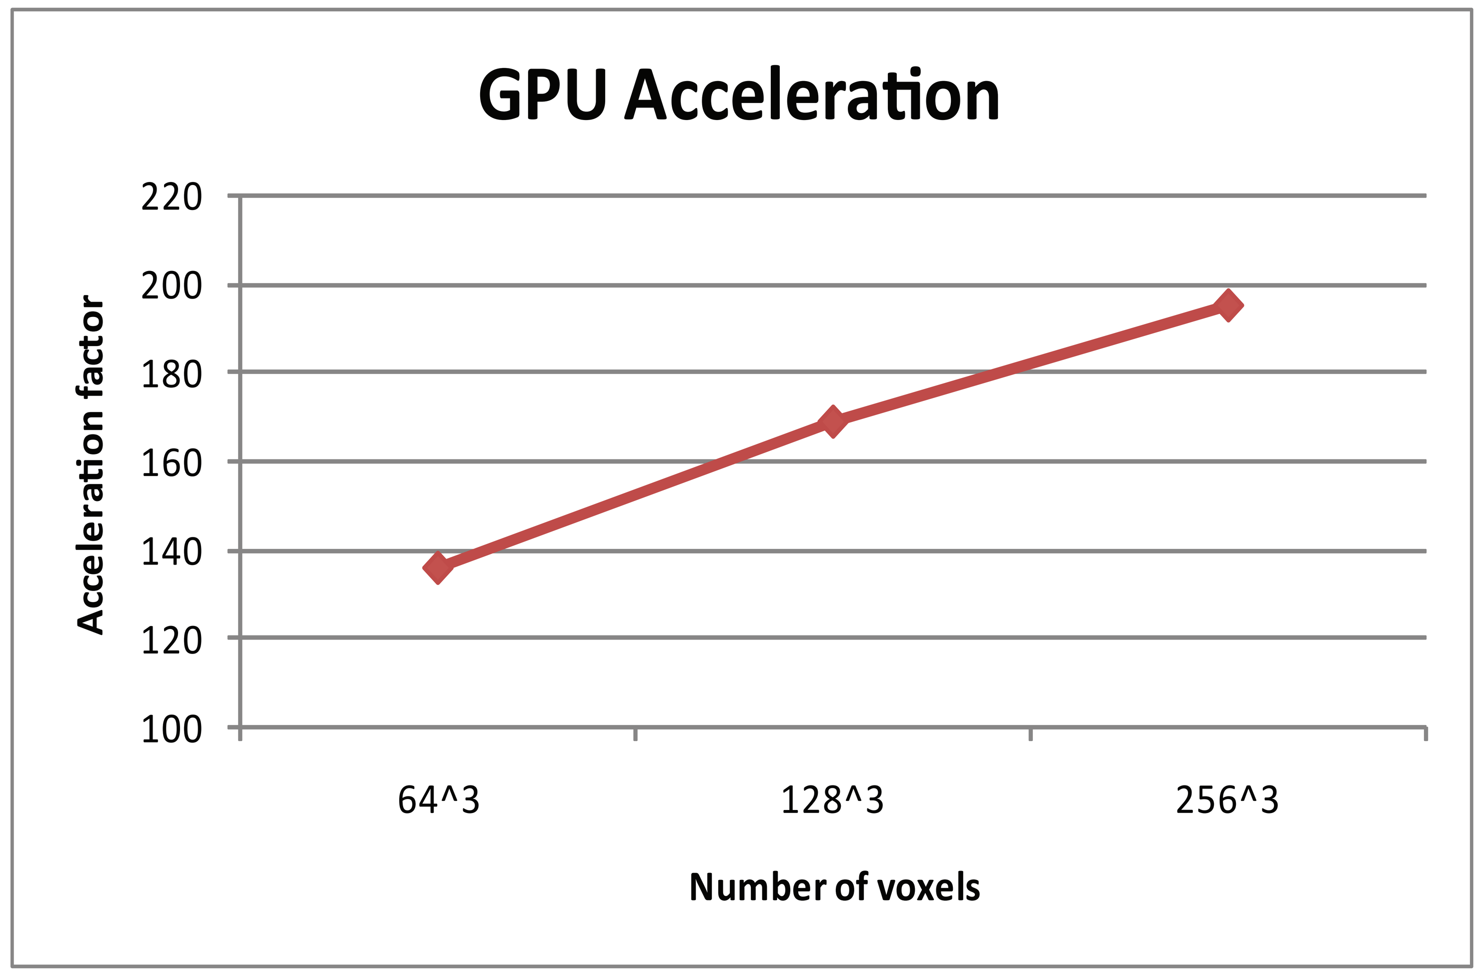

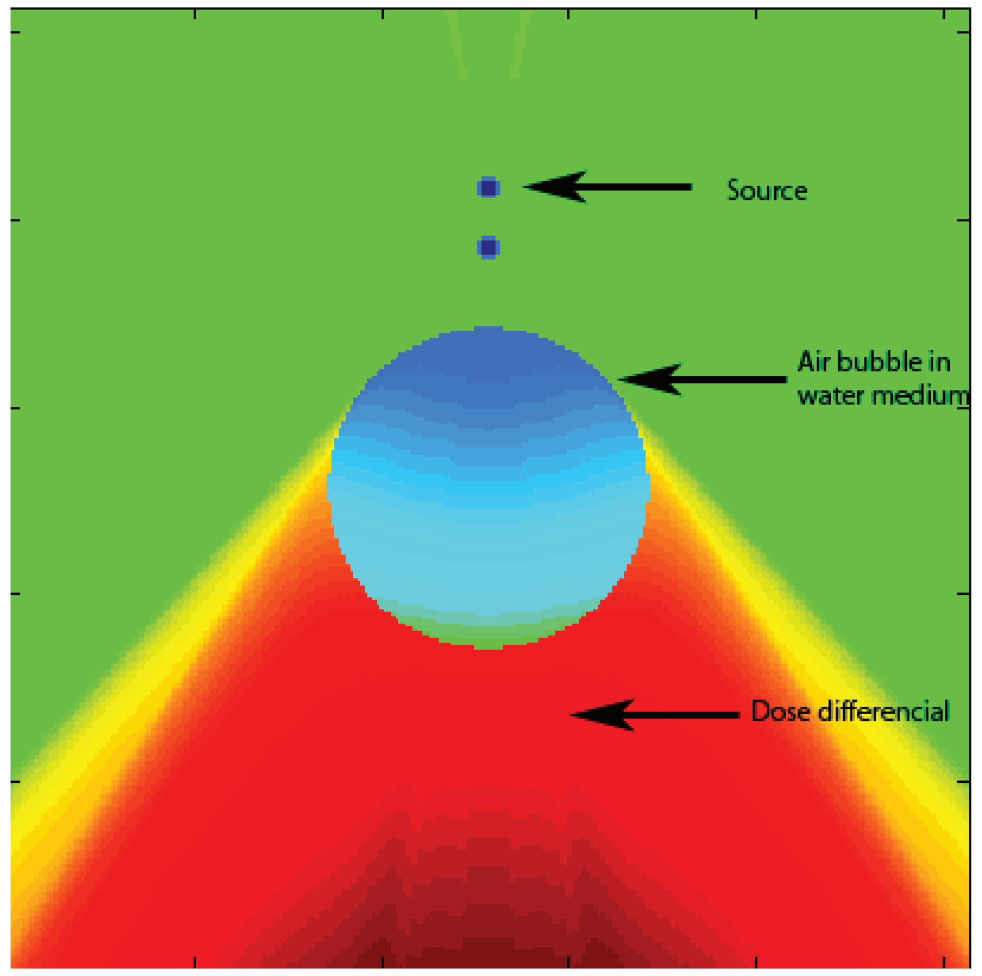

Conclusion

We have studied the effects of using graphics hardware to compute the dose in radiotherapy treatment planning.
Our acceleration factors, ranging from 30x to 930x, coupled with an equal or enhanced physical accuracy when compared to our single threaded CPU implementation show the tremendous potentiel of the platform.
Future projects include integrating the dose computation modules in an Intensity Modulated Ratiotherapy Treatment (IMRT) optimization where the dose has to be computed several thousand times.