PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

A real-time radiation dose monitoring system for patients and staff during interventional fluoroscopy using a GPU-accelerated Monte Carlo simulator and an automatic

Andreu Badal, Fahad Zafar, Han Dong, Aldo Badano

Andreu Badal, Fahad Zafar, Han Dong, Aldo Badano, "A real-time radiation dose monitoring system for patients and staff during interventional fluoroscopy using a GPU-accelerated Monte Carlo simulator and an automatic 3D localization system based on a depth camera," Proc. SPIE 8668, Medical Imaging 2013: Physics of Medical Imaging, 866828 (19 March 2013); doi: 10.1117/12.2008031

SPIE.

Event: SPIE Medical Imaging, 2013, Lake Buena Vista (Orlando Area), Florida, United States

A real-time radiation dose monitoring system for patients and staff during interventional fluoroscopy using a GPU-accelerated Monte Carlo simulator and an automatic 3D localization system based on a depth camera

Andreu Badal[a,*], Fahad Zafar[a,b], Han Dong[a,b] and Aldo Badano[a]

[a] Division of Imaging and Applied Mathematics, OSEL, CDRH, U.S. Food and Drug Administration, 10903 New Hampshire avenue, Silver Spring MD 20993, USA;
[b] Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore MD 21250, USA

ABSTRACT

Radiation monitoring systems able to accurately track the radiation dose received by the patient and the medical staff during interventional fluoroscopy can be used to minimize the likelihood and severity of radiation-induced skin injuries and estimate the accumulated organ doses. We describe a method to monitor doses in real time using automatic sensors in the imaging room and a GPU-accelerated computer simulator. The Monte Carlo simulation code MC-GPU is used to estimate patient and staff doses due to primary and scattered radiation, along with the associated statistical uncertainties. The geometrical configuration of the irradiation is automatically determined and updated using data from a depth camera that tracks the location and posture of each person in the imaging room. A virtual x-ray source graphical interface is used to manually trigger the simulations. The implemented computational framework separates the simulation of the x-ray transport through the patient and the operator bodies into two coupled, sequential simulations. The initial simulation uses the patient anatomy and a c-arm source model with a collimated cone beam emitted from a point focal spot. During this simulation a large phase space file with the energy, position and direction of x rays scattered in the direction of the operator is created. The phase space file is then used as the input radiation source for the following simulation with the operator anatomy model. Particle recycling is employed as a variance reduction technique to maximize the information obtained from the limited number of particles scattered towards the operator. For a typical image acquisition, a patient skin dose map can be displayed at the operator's monitor within 10 seconds with a peak skin dose error below 1%. This work demonstrates that a dose monitoring system based on accurate Monte Carlo simulations can be used to estimate in real-time the average and peak organ doses for both the patient and the staff in interventional fluoroscopy, and provide timely information regarding possible overdoses while the imaging procedure is being performed.

Keywords: Monte Carlo, dosimetry, GPU, depth camera, fluoroscopy

1. INTRODUCTION

In most x-ray imaging procedures patients are exposed to low levels of radiation such that radiation injuries are rarely encountered. An exception is interventional fluoroscopy, a modality in which it may be clinically acceptable to deliver high doses of radiation to patients that are in critical condition (for example, patients with an obstructed coronary artery). Severe radiation injuries, especially in the skin, after fluoroscopic interventions have been extensively reported in the literature.[1–4] It is considered a good practice to follow up patients that may have been exposed to high doses (for example, a peak skin dose greater than 2 Gy) to make sure that they get the appropriate treatment in the event of developing a radiation injury, which may take weeks or months to manifest.[3]

* Corresponding author: Andreu.Badal-Soler@fda.hhs.gov

Medical Imaging 2013: Physics of Medical Imaging, edited by Robert M. Nishikawa, Bruce R. Whiting, Christoph Hoeschen,
Proc. of SPIE Vol. 8668, 866828 · © 2013 SPIE · CCC code: 1605-7422/13/$18 · doi: 10.1117/12.2008031

An important difference between interventional fluoroscopy and other radiographic modalities is that the radiologist or other physician performing the intervention and auxiliary medical personnel may stay inside the imaging room, close to the patient, during the x-ray exposure and therefore they are exposed to radiation scattered from the body of the patient. The occupational radiation exposure of the medical personnel working in the imaging room is much lower than the patient dose but it is important to monitor the dose to ensure that it does not exceed recommended maximum annual levels.[5]

In this work a method to monitor skin and organ doses in real time using automatic sensors in the imaging room and a computer simulator in a remote location is presented. The purpose of the dose monitoring system is to inform the operators of the distribution of the dose received by the patient's skin and its peak value, and keep a record of approximate organ doses received during the procedure that can be used to track the total doses received over time. The system can also be used for educational purposes to inform operators of the safest locations in the room to work and display a warning message when the calculated dose to the operator is higher than a certain threshold value. Other dose tracking systems have been developed by different groups.[6,7] A distinctive difference of our system is the use of detailed, state-of-the-art particle transport models for dose estimation instead of less accurate methods. The accurate modeling of scattered radiation in the presented system allows for monitoring the dose to the operator of the x-ray equipment, which is not measured in conjunction with the patient dose in any other system.

2. METHODS

A computational framework for dose monitoring in interventional fluoroscopy is presented in the following subsections. At this stage of development, the system is not meant for clinical use and it represents a proof of concept to investigate the limits of some new technologies with a promising future in the field of computational dosimetry.

The two main hardware components of the framework are a control computer and a depth camera, both located inside the imaging room. A Monte Carlo (MC) simulation code (see section 2.1) running in the control computer or, optionally, on a remote computer cluster (see section 2.3) is used to estimate the 3D dose distribution on the skin and the internal organs of both the patient and the medical staff. The geometrical configuration of the irradiation is automatically determined using data from the depth camera, which continuously tracks the location and posture[8] of each person in the imaging room (see section 2.2). The computed radiation doses can be displayed at the control computer as soon as they are available. A detailed description of the implemented algorithms is given in section 2.4.

2.1 Monte Carlo Simulations and Anatomical Models

The Graphics Processing Unit (GPU)-accelerated MC simulation code MC-GPU[9,10] is used as the computational engine to model x-ray transport and estimate the radiation dose to the patient and the operator of the imaging equipment. MC-GPU implements the x-ray transport physics models from PENELOPE 2006,[11] which have been exhaustively validated in the past.[12,13] MC-GPU employs a voxelized geometry model with an efficient particle transport algorithm based on delta scattering (Woodcock tracking). The interaction sampling and geometry ray-tracing algorithms were designed to provide an optimum performance in GPUs, minimizing the accesses to the slow video memory while maximizing the parts of the code that can be executed in parallel in thousands of concurrent GPU threads. The simulation code can run in parallel in multiple GPUs in multiple computers, as described in section 2.3.

The detailed anatomical models from the Virtual Family[14] were used to represent the internal human anatomy in the simulations. This set of computational phantoms includes data for four possible patients: a 34-year-old adult male, 26-year-old female, and 11- and 6-year-old children. We used rigid human phantoms in supine position, but the data from a 3D camera provides information on the location of the legs, arms and head of each person that could be potentially used to modify the phantom's pose in the future. The user of the code can select in the input file the phantom to be used in the simulation and a scaling factor for each dimension of the phantom. With this scaling factor the model can be adapted to a particular body size. However this simple scaling method is only adequate for small modifications of patient size because it does not correctly model the change in organ size for different body sizes (for example, taller people do not necessarily have taller hearts).

The radiation shielding protections that are commonly used during fluoroscopic imaging (including leaded aprons, eyeglasses or drapes) and the patient bed are not currently modeled by our software, but many of these elements could be easily included in the voxelized models. Since the radiation shields are not modeled, the system will provide an estimate of the maximum dose that could be received in a worst-case scenario irradiation. It is important to mention that the use of a standard human anatomy model instead of a personalized model is a major source of error in our system, and a fact that will inevitably prevent the system from computing the exact real organ doses independently of the accuracy of the MC code. The magnitude of the sources of error affecting the simulations should be evaluated in detail during the system validation process before the system can be tested in a clinical setting.

As mentioned before, MC-GPU can be executed in parallel in multiple computers. The simulations presented in this work were executed on the DIAM GPU cluster containing 14 GPUs: 8 NVIDIA GeForce GTX 580 and 6 NVIDIA GeForce GTX 680. The cluster nodes had a 6-core Intel Xeon CPU (2.40 GHz) running the Linux operating system (Ubuntu 12.04), the NVIDIA CUDA 5.0 library and the NVIDIA GPU driver version 304.43.

2.2 Depth Cameras

In order to compute the dose to the operator of the fluoroscopic equipment, a means to determine the exact position of the operator at the moment of each irradiation is needed. We decided to use a consumer-grade depth camera to track the medical personnel inside the imaging room. The data from the 3D camera is used to automatically configure the simulation geometry. Two similar cameras were evaluated in this work: a Kinect for Xbox 360 (Microsoft Corp., Redmond WA) and an ASUS XtionPRO (ASUSTEK Computer Inc., Taipei, Taiwan)*. Both cameras use the depth sensor and processing chip developed by PrimeSense LTD (Tel-Aviv, Israel). This depth-sensing technology is based on the use of an infrared laser and a holographic filter to project a known pattern towards the room. A video camera with an optical filter to block visible light is used to record the scene in the infrared range. A microprocessor chip in the camera post-processes the infrared image to estimate the distance of the objects in front of each pixel. To be accurately detected, the objects must be located at a distance of between 0.5 and 6 meters approximately. The depth data is transmitted to the computer as a 640×480 pixel image where each pixel value corresponds to the distance of the object behind that pixel. The depth value is computed with 11-bit precision (2048 levels) but it can be observed that the depth estimations have a noticeable level of noise (jitter) of the order of ± 1 mm. The main advantages of these two consumer-grade cameras are the price ($\sim$\$150) and the fact that there is a large community of developers creating freely-available libraries and applications that simplify the programming of new tools. Even though both cameras could be used in the described framework, we used the ASUS camera instead of the Kinect because of the smaller physical size, the available USB connector and the lack of an external power supply requirement.

The open source OpenNI framework (`http://www.openni.org/`) was used in this work to operate the sensor from a personal computer and to have access to the raw depth data. The computer vision middleware NiTE™ (Natural Interface Technology for End-User, `http://www.openni.org/files/nite/`) was used to process the scene and separate the human bodies from the background. This software is also able to detect and track a skeleton joint map of the person which makes it simple to estimate the pose and orientation of the body. Other software libraries and applications exist to extract further information from the depth data, such as facial expressions or hand movements, but these tools were not required in our project.

2.3 Interprocess Communication Protocols

The computational framework uses three separate computer applications to deal with different aspects of the image acquisition process. These three applications need to share information and synchronize their execution. For example, the MC simulations can not start until information about the x-ray source parameters and the operator location is available. The communications between these three applications, which all run in the local workstation at the imaging site, are handled using the open source middleware system Lightweight Communications and Marshalling (LCM, `https://code.google.com/p/lcm/`). LCM is a set of libraries and tools for

*The mention of commercial products, their sources, or their use in connection with material reported herein is not to be construed as either an actual or implied endorsement of such products by the Department of Health and Human Services.

inter-process communication via UDP (User Datagram Protocol) message passing specifically designed for real-time systems where high-bandwidth and low latency are critical. The main advantages of this library are the simplicity of the message passing functions and the multi-language compatibility that allows the communication of applications written in C, C++, Java, Python and MATLAB.

Apart from the LCM communications, the MC simulator also uses the Message Passing Interface library (MPI, `http://www.open-mpi.org/`) to access multiple GPUs in different computers in a remote computing cluster. MPI uses secure shell connections to launch computational threads in the remote computers and can access any system available in the local network or on the internet. The latency of the secure connections is not negligible and therefore the control computer will have to wait some time to receive the final results from all the remote nodes. With the Gigabit ethernet connections in our intranet the total time spent transferring the large dose maps across the network accounted for less than 5 seconds of overhead in typical simulations.

2.4 Implemented Algorithm for Radiation Dose Monitoring

The implemented dose monitoring framework consists of three separate computer codes: the MC simulator module, the virtual x-ray source console module and the depth camera tracking module. Figure 1 presents a detailed flow chart describing the different elements of the software framework.

The acquisition technique factors and other relevant imaging parameters (kVp, filtration, collimation, gantry position and angulation) are manually input by the system operator using the virtual x-ray source console graphical user interface (GUI). A picture of the source GUI is shown in Fig. 2. To start an image acquisition the user just needs to click a button in the GUI. At the moment the button is clicked the image acquisition parameters are broadcasted to all applications running in the workstation with an LCM message. These messages trigger the acquisition of the operator position in the depth camera module. The operator position is then broadcasted again with LCM messages. The MC simulation module waits for the messages from the source and the camera to start the simulation (synchronous message reception). In a clinical implementation, the simulations could be

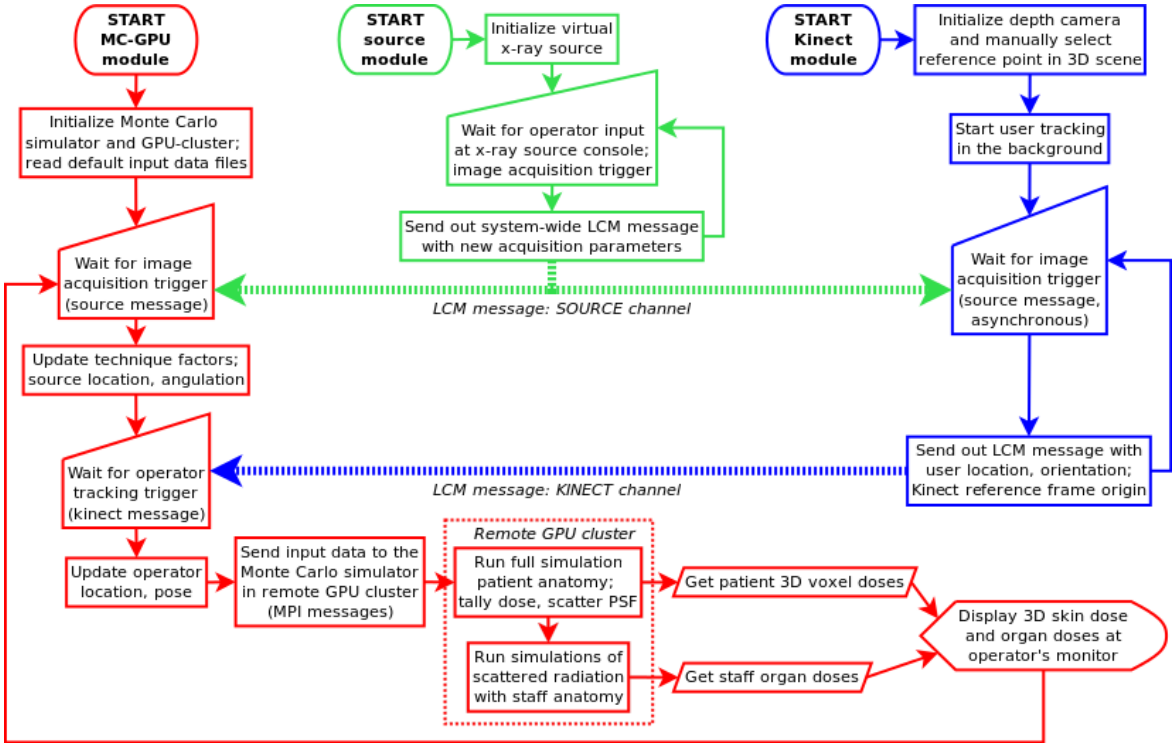

Figure 1. Schematic flow diagram of the proposed dose monitoring computational framework composed of three separate applications (x-ray source module, Kinect module, MC-GPU simulator module) communicating through LCM messages.

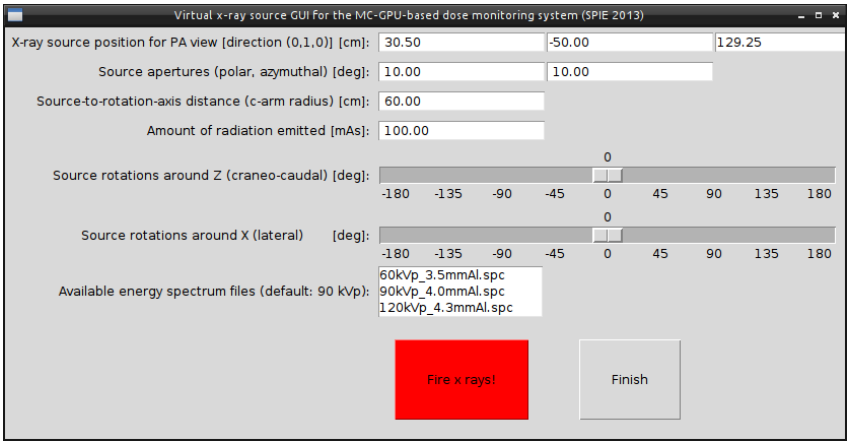

Figure 2. Graphical user interface of the virtual x-ray source.

triggered by an external radiation monitoring device such as a dose-area product (DAP) sensor instead of by messages from the x-ray console. The measured DAP value could also be used to convert the simulated relative doses (eV/g per x-ray) to the appropriate absolute dose units (Gy) using pre-calculated conversion factors.

The depth camera tracking software runs during the entire imaging procedure. A live video stream showing the scene depth map and the operator skeleton tracking marks is displayed during the program execution. An example depth image is shown in Fig. 3. The camera module is asynchronously subscribed to the x-ray source LCM messaging channel to receive a notice whenever an exposure takes place without interrupting the scene tracking. In response to the incoming message from the source, the camera module sends out an LCM message with the location and orientation of the operator body at the time of the acquisition (asynchronous message reception). Before the first image acquisition takes place, the user must find and click in the depth image the location of three landmarks that will define the reference system used in the simulation with the patient. First, the user has to click at the location that corresponds to the origin of coordinates of the simulation geometry. By default, the origin is located at the lower back corner of the patient bounding box, i.e., a point below the left feet of the patient. Then, the user has to click a point above the origin to define the patient posterior-anterior direction, and a point towards the head to define the caudal-cranial direction. The remaining left-right direction is defined by the cross-product of the two other dimensions.

The MC simulation module is divided in two separate parts: an initialization part that is executed before the medical procedure begins, and a dosimetric simulation loop that is executed for each irradiation. During the initialization part, the large voxelized geometry and material files are read from disk and the interaction sampling models are initialized. The MPI and GPU contexts are also initialized and the required simulation tables and constant parameters are transferred to the slow GPU memory. The dosimetric simulation part starts at the beginning of each image acquisition, triggered by the signal from the source module. The simulator fetches the technique factor data from the source messages and the location of each person in the room from the depth camera messages. Then, two simulations are successively launched in the local workstation or the remote cluster: a short simulation ($\sim$5 seconds) with the patient anatomy to compute the patient average and peak organ doses and 3D dose map, and a separate simulation with the operator anatomy to estimate the operator doses. The dosimetric results are reported at the local workstation as soon as all the computational nodes finish their simulations. The dose deposited in each organ can be readily visualized with a provided set of GNUPLOT scripts. The complete 3D dose distributions can also be visualized (for example using IMAGEJ) to provide more information to the operator.

The key innovation that allows the system to compute the operator dose is that during the simulation with the patient anatomy all the particles escaping the simulation universe are tested for intersection with the operator bounding box. Those x-rays scattered in the direction of the operator are stored in a phase space file (PSF) in GPU memory. When the simulation with the patient finishes, a new simulation with the operator anatomy is

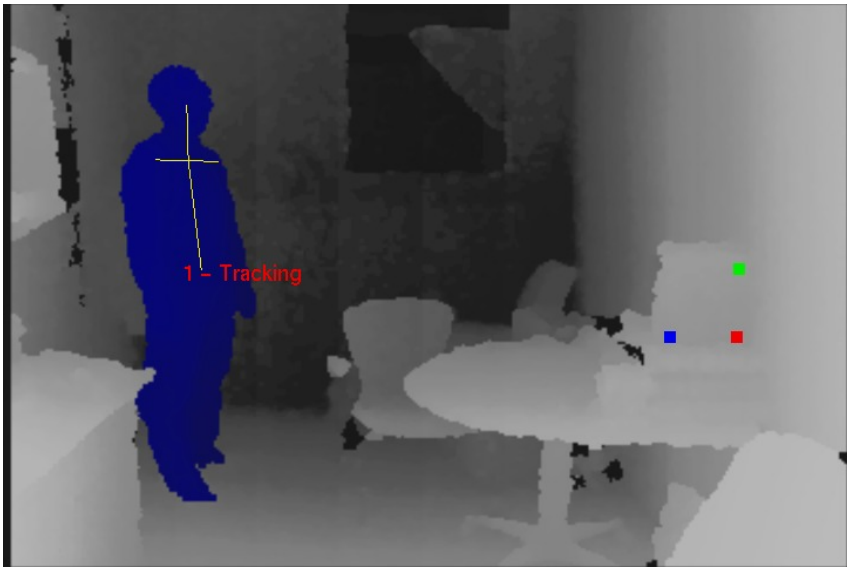

Figure 3. Visualization of a room with a person provided by the depth camera tracking module. Four landmarks are automatically detected on the body: right shoulder, left shoulder, head center and hip center. The three color dots are user-selected landmarks that define the patient reference system in the posterior-anterior and caudal-cranial directions.

started using the PSF as the x-ray source. Since only a small fraction of the initial x rays are expected to be scattered towards the operator, each x ray in the PSF is recycled 128 times to maximize the information obtained from each particle (variance reduction). The dose deposited by each x ray is scaled down a factor 128 to keep the mean deposited dose estimation unbiased. The amount of particles in the PSF depends strongly on the distance from the operator to the beam entrance point (inverse square law). As a consequence, the simulation time for the operator will be longer, and the accuracy of the results better, when the operator is located near the source.

It is important to mention that in a clinical situation fluoroscopy or cine-angiography images may be continuously acquired for multiple seconds or minutes. For the sake of simplicity we compute the doses only at the end of each sequence. The dose monitoring system determines the position of the operator and the source angulation at the beginning of each irradiation sequence and it is assumed that the staff, equipment or patient do not move during the sequence. A more complex simulation scheme with multiple time steps could be implemented if sufficient computational resources were available. Off-focus radiation leaking from the source or scattered from the beam collimators is assumed to be negligible and not considered in our model.

2.5 Sample Coronary Angiography Simulations

The presented dose monitoring system was tested simulating two idealized coronary angiography procedures. The adult male phantom "Duke" from the Virtual Family was used to represent the patient and the operator anatomy in the simulations. This voxelized phantom is composed of 9 different materials and has a size of $61 \times 31 \times 186$ cm^3, with 0.2 cm voxels (350 MBytes of GPU memory). A 90 kVp energy spectrum with 4.0 mm Al filtration was used for all exposures. The simulations were executed in parallel in 14 GPUs. It was assumed that all the exposures had the same amount of radiation (same mAs or DAP) and radiation protection shields were not used.

In the first simulated case, four posterior-anterior fluoroscopic projections were acquired with the operator moving 1 m between exposures. In the first exposure the operator was located at the level of the patient left knee, close to the bed; in the second he was located at the level of the chest, close to the x-ray beam; in the third he moved 1 m away from the knee level; and in the last exposure he was located 1 m away from the chest (see Fig. 4, above, for a graphical representation). The purpose of this study was to investigate the effect of the operator position on the received dose.

In the second case, twelve fluoroscopic projections were simulated with different cranial and lateral angulations separated by 45˚ increments. In this case the operator remained static at the level of the patient knee, close to the bed. The purpose of this study was to investigate how the dose received by the patient and the operator change depending on the c-arm angulation.

3. RESULTS

3.1 Case 1: dose for different operator locations

Figure 4 shows volume renderings of the 3D dose distributions received by the operator for each of the four irradiation events in the first example coronary angiography simulation described above. The statistical uncertainty in the patient voxel dose estimation at the beam entrance region (including the peak skin dose) was below 1%. The average organ doses for all organs inside or near the field of view were also estimated within 1% uncertainty (note that estimating the average dose to complete organs requires much fewer histories than estimating the peak dose tallying the dose to individual voxels). The part of the body that received the largest amount of radiation were the legs, because the x-ray source was located below the patient's bed and most scattered radiation was generated at the back of the patient. In clinical facilities it is common to install leaded drapes at the sides of the bed to block this source of scattered radiation. As expected, our results suggest that blocking the radiation coming from below the bed is a good method to reduce the dose to the operator in a posterior-anterior projection. The results also demonstrate that increasing the distance between the operator and the x-ray source is an effective method to reduce the radiation exposure.

The presented simulations were executed in parallel in 14 GPUs with an approximate simulation time per projections of 5 s. The average speed of the simulation in one NVIDIA GeForce GTX 580 GPU was $10.4{\cdot}10^6$ x-ray/s. The total number of x rays simulated in all GPUs for each projection was $7.3{\cdot}10^8$. Figure 5 shows a plot of the time spent simulating each projection, separating the time spent in the patient simulation and in the

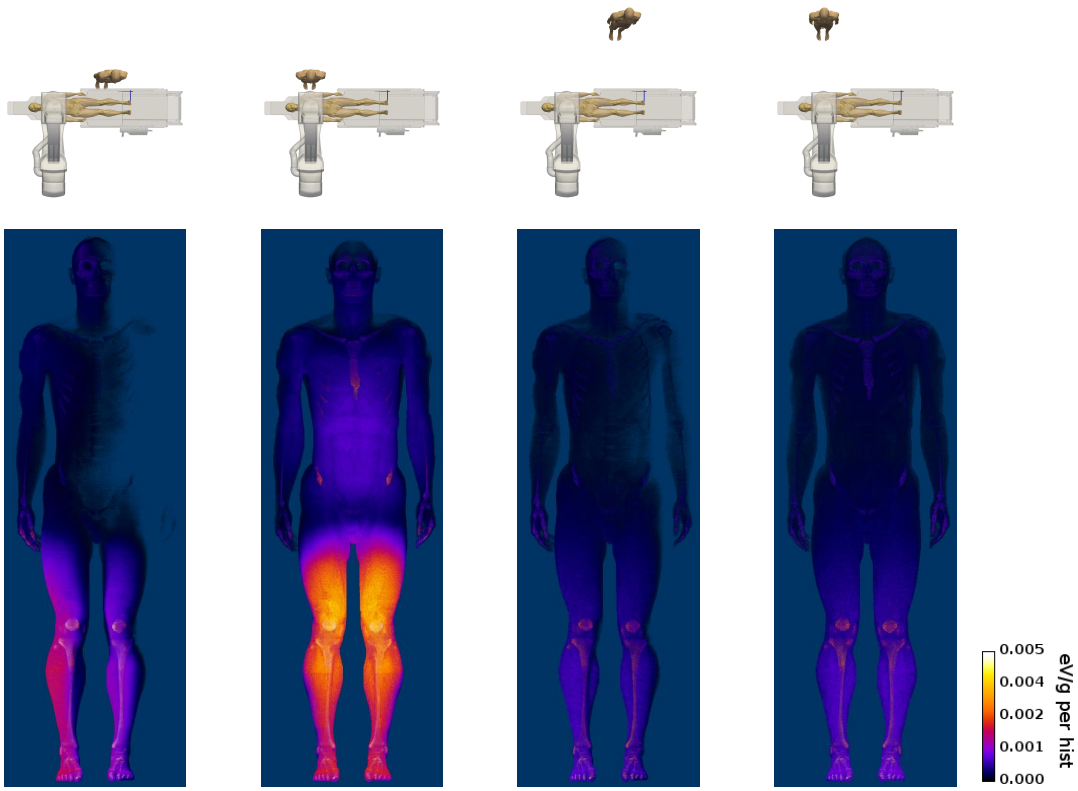

Figure 4. Results of the simulated case 1 with the operator at four different positions separated by 1 m: visualization of the imaging room setting (above), volume rendering of the estimated 3D dose distributions to the operator (below).

operator simulation. The figure also reports the total energy absorbed in the patient and the operator for each projection (the code reports the mean and peak doses for each organ but this information is not plotted). The dose to the patient was constant in this case because the source angulation was fixed, but the operator dose was reduced 3 times when the operator moved 1 m away from the patient chest region. For the complete program execution, the total time spent in the MC transport was 41.2 s and the time spent in initialization and reporting was 199.2 s. Therefore on average the simulation required 10 s per projection. However, the simulation time for the operator had a wide variability due to the fact that the number of particles in the PSF was much larger when the operator was close to the x-ray beam. The number of particles in the PSF for each projection was $10.3{\cdot}10^6$, $43.6{\cdot}10^6$, $3.6{\cdot}10^6$ and $5.3{\cdot}10^6$ respectively. A factor 10 increase in the number of particles in the PSF for the second projection (operator closest to the beam) corresponded to a factor 10 increase in the simulation time, as expected.

The average and peak skin doses deposited in the patient for each projection ($\pm 2\sigma$) were 0.2243 ± 0.0001 and 44.93 ± 0.02 eV/g per history. When the operator was located near the chest of the patient, the operator average and peak skin doses were 0.00063 and 0.00444 eV/g per history. When the operator moved 1 m away from the chest, the operator doses were 0.00021 and 0.00168 eV/g per history. Therefore the average patient skin dose was 356 larger than the operator skin dose in the first case, and 1068 times larger in the second. The patient peak skin dose was 10243 times larger than the operator peak skin dose in the first case, and 26969 times larger in the second. A difference of three to five orders of magnitude between the skin doses received by the patient (exposed to the primary x-ray beam) and the operator (exposed to scattered radiation) is logical and explains why deterministic radiation effects have not been observed in operators.

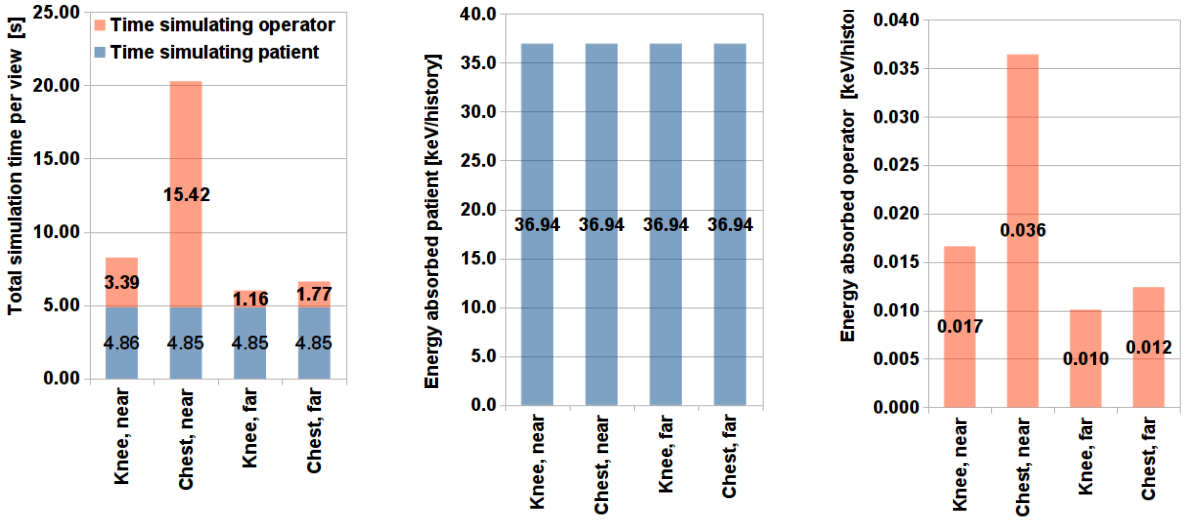

Figure 5. Time spent simulating the patient and operator doses for the four views, and total energy absorbed in the patient and operator bodies. Uncertainties below 1% (dose to the patient constant because source angulation was fixed).

3.2 Case 2: dose for different c-arm angulations

Figure 6 shows a volume rendering of the 3D dose distribution received by the patient after the twelve irradiation events described in the second example coronary angiography case (see section 2.5). It can be observed that the peak skin dose was located at the back of the patient, at the region where three different projections overlapped. The statistical uncertainty in the estimated patient peak skin dose and average organ doses was below 1%.

Figure 7 shows five volume renderings of the 3D dose distributions received by the operator in five of the twelve projections. The maximum dose to the operator was received for the anterior-posterior projection (x-ray source above the patient, second image in the figure). Figure 8 presents the simulation time for each projection and the total energy absorbed in the patient and the operator. It was observed that the patient peak skin dose (not shown) varied much more than the total energy absorbed in the whole body for different angulations.

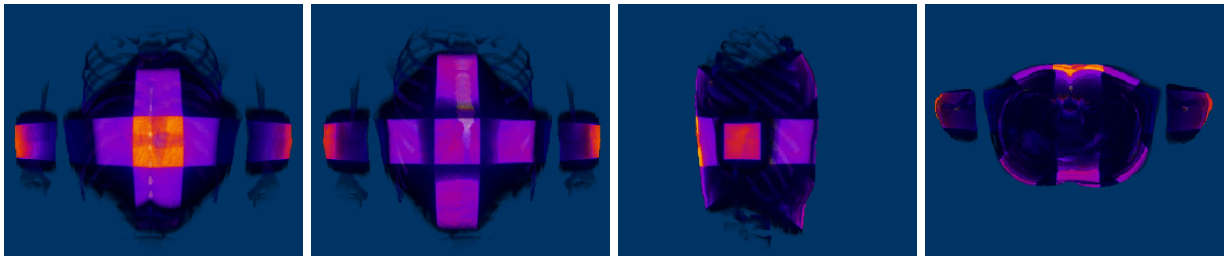

Figure 6. Volume rendering of the patient 3D dose distribution after twelve x-ray exposures at 45° increments. The images (from left to right) were rendered from the back of the patient, the front, the right side and from above.

Equivalently, the operator dose also varied significantly depending on the location of the x-ray source. It is worth mentioning that the dose estimates would be different if we had modeled an automatic exposure control system instead of assuming that all projections had the same intensity. For example, our results show that the dose is practically null for the lateral projection with the source on the opposite side of the operator. Using automatic exposure control, the intensity of the beam (and the resulting doses) for the lateral projections would be significantly increased in order to obtain a usable projection image. In this case our results would still correctly predict that the dose to the operator would be larger if the operator was located at the side of the x-ray source compared to staying at the opposite side, next to the detector.

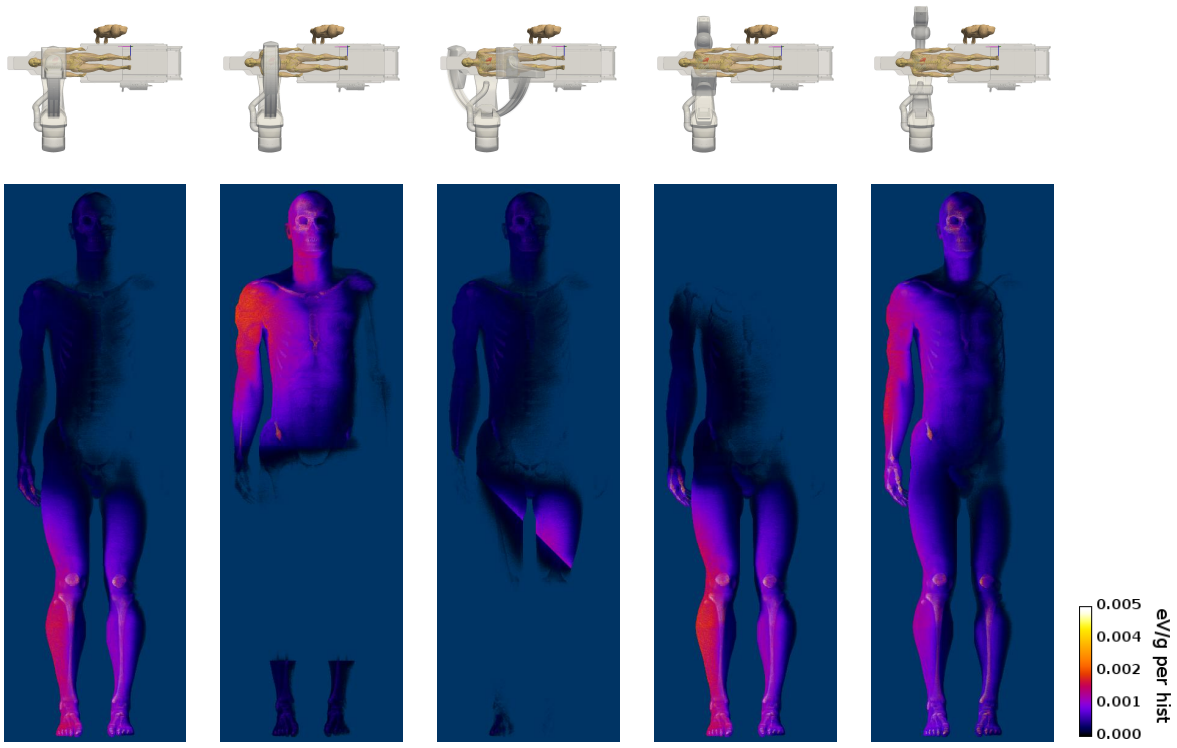

Figure 7. Results of the simulated case 2 with 12 different c-arm angulations: visualization of 5 sample imaging room settings (above), volume rendering of the estimated operator 3D dose distributions for posterior-anterior, anterior-posterior, 45° caudal, 45° right and 90° right (lateral) angulations (below).

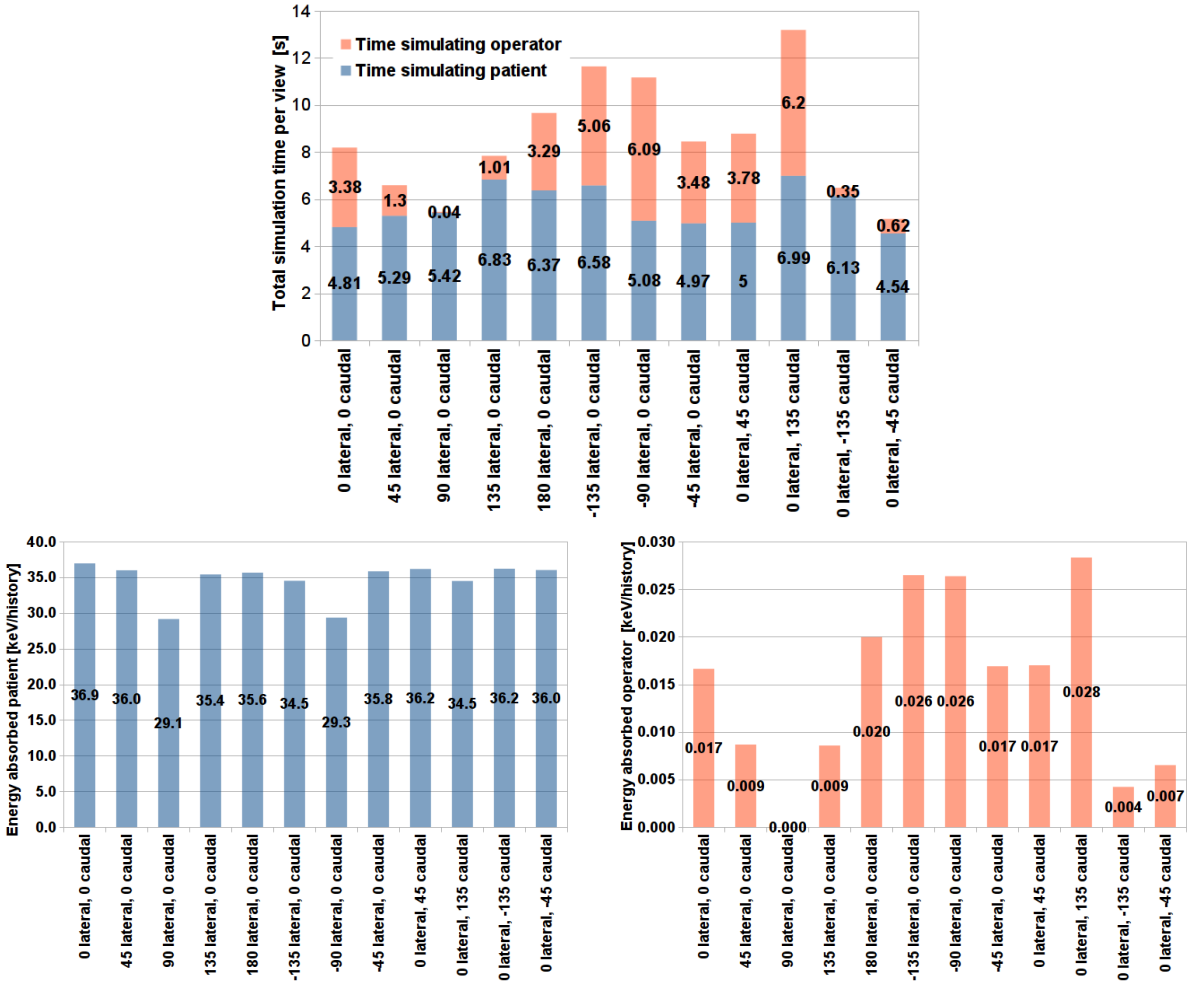

Figure 8. Time spent simulating the patient and operator doses for the twelve views (above), and total energy absorbed in the patient and operator bodies (below). Uncertainty in the absorbed energy estimations below 1%.

4. CONCLUSIONS

A dose monitoring system that uses an accurate MC code, detailed anatomical phantoms and physical sensors in the imaging room has been presented. The system has the potential to provide real-time dose estimations for both patients and staff during interventional fluoroscopy with high accuracy. This work demonstrates that GPU-accelerated MC simulation is a valid alternative to less accurate methods for real-time dosimetric applications.

Work is underway to improve the dose tracking system and fix some of the limitations of the current software, for example: model the radiation shields, implement a method to convert the dosimetric results to Gy, validate the accuracy of the system by comparing with experimental measurements, and estimate the magnitude of the different sources of error in the dose estimations.

ACKNOWLEDGMENTS

We would like to acknowledge the contribution of Dr. Donald Miller to this work, with useful discussions and inspiration for the initial design of the project. The simulations presented in this work could not have been possible without the computational resources available at the Division of Imaging and Applied Mathematics (OSEL, CDRH, FDA) and the assistance of the system administrator Jonathan Boswell.

REFERENCES

[1] Shope, T. B., "Radiation-induced skin injuries from fluoroscopy.," *Radiographics* **16**, 1195–1199 (1996).

[2] Miller, D., Balter, S., Cole, P., Lu, H., Schueler, B., Geisinger, M., Berenstein, A., Albert, R., Georgia, J., Noonan, P., Cardella, J., George, J. S., Russell, E., Malisch, T., Vogelzang, R., 3rd, G. M., and Anderson, J., "Radiation doses in interventional radiology procedures: the RAD-IR study. Part I: overall measures of dose," *J Vasc Interv Radiol* **14**, 711–727 (2003).

[3] Miller, D., Balter, S., Cole, P., Lu, H., Berenstein, A., Albert, R., Schueler, B., Georgia, J., Noonan, P., Russell, E., Malisch, T., Vogelzang, R., Geisinger, M., Cardella, J., George, J. S., 3rd, G. M., and Anderson, J., "Radiation doses in interventional radiology procedures: the RAD-IR study. Part II: skin dose," *J Vasc Interv Radiol* **14**, 977–990 (2003).

[4] Balter, S., Hopewell, J. W., Miller, D. L., Wagner, L. K., and Zelefsky, M. J., "Fluoroscopically guided interventional procedures: a review of radiation effects on patients' skin and hair," *Radiology* **254**, 326–341 (2010).

[5] NCRP Report No. 168, "Radiation dose management for fluoroscopically guided interventional medical procedures," Tech. Rep. 168, National Council on Radiation Protection and Measurements (2010).

[6] Bednarek, D. R., Barbarits, J., Rana, V., Nagaraja, S., Josan, M., and Rudin, S., "Verification of the performance accuracy of a real-time skin-dose tracking system for interventional fluoroscopic procedures," in [*Medical Imaging 2011: Physics of Medical Imaging*], *Proc. SPIE* **7961**, 796127 (2011).

[7] Johnson, P. B., Borrego, D., Balter, S., Johnson, K., Siragusa, D., and Bolch, W. E., "Skin dose mapping for fluoroscopically guided interventions," *Med. Phys.* **38**, 5490–5499 (2011).

[8] Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A., and Blake, A., "Real-time human pose recognition in parts from single depth images," *CVPR* **3** (2011).

[9] Badal, A. and Badano, A., "Accelerating Monte Carlo simulations of photon transport in a voxelized geometry using a massively parallel Graphics Processing Unit," *Med. Phys.* **36**, 4878–4880 (2009).

[10] Badal, A., Kyprianou, I., Sharma, D., and Badano, A., "Fast cardiac CT simulation using a Graphics Processing Unit-accelerated Monte Carlo code," in [*Medical Imaging 2010: Physics of Medical Imaging*], Samei, E. and Pelc, N. J., eds., *Proc. SPIE* **7622**, 762231–762231–9 (2010).

[11] Salvat, F., Fernández-Varea, J. M., and Sempau, J., [*PENELOPE - A code system for Monte Carlo simulation of electron and photon transport*], Nuclear Energy Agency (OECD) Issy-les-Moulineaux (2006). Available in pdf format at `www.nea.fr`.

[12] Sempau, J., Fernández-Varea, J. M., Acosta, E., and Salvat, F., "Experimental benchmarks of the Monte Carlo code PENELOPE," *Nucl. Instrum. Meth. Phys. Res. B* **207**, 107 – 123 (2003).

[13] Ye, S.-J., Brezovich, I. A., Pareek, P., and Naqvi, S. A., "Benchmark of PENELOPE code for low-energy photon transport: dose comparisons with MCNP4 and EGS4," *Phys. Med. Biol.* **49**, 387 – 397 (2004).

[14] Christ, A., Kainz, W., Hahn, E. G., Honegger, K., Zefferer, M., Neufeld, E., Rascher, W., Janka, R., Bautz, W., Chen, J., Kiefer, B., Schmitt, P., Hollenbach, H. P., Shen, J. X., Oberle, M., Szczerba, D., Kam, A., Guag, J. W., and Kuster, N., "The Virtual Family–development of surface-based anatomical models of two adults and two children for dosimetric simulations," *Phys. Med. Biol.* **55**, N23 – N38 (2010). Models available at `http://www.itis.ethz.ch/services/anatomical-models/`.