Fast Monte Carlo proton treatment plan validation in the Google Cloud

Andrew F Green[1,2], Adam H Aitkenhead[3], Hywel L Owen[1,2] and Ranald I Mackay[3]

[1] School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK
[2] Cockcroft Institute, Daresbury Science and Innovation Campus, Daresbury WA4 4AD, UK
[3] The University of Manchester, Manchester Academic Health Science Centre, The Christie NHS Foundation Trust, Wilmslow Road, Manchester M20 4BX

E-mail: hywel.owen@manchester.ac.uk

Abstract. Monte Carlo validation of proton therapy treatment plans is one method to improve the quality of delivered irradiation by reducing the dose calculation uncertainty, but its use is hampered by its typically slow speed in comparison to the overall planning and treatment workflow. Most attempts at fast Monte Carlo calculation have hitherto used locally-based clusters or graphical processing units, usually combined with a simplified implementation of the proton transport. Here we present the use of cloud computing to perform rapid Monte Carlo dose estimation using the comprehensive physics model implemented within GEANT4. Mock proton treatment plans in a head-and-neck phantom were validated using 1600 Google Cloud CPUs. The plan consisted of 16 million voxels each approximately $1 \times 1 \times 2$ mm^3 in size, and 10 million primary particles were tracked, giving a dose uncertainty of 1% within the high-dose region. The validation was completed in just 6 minutes and cost \$11.20 (£7.17); reduced execution time is possible with more CPUs, bringing real-time validation within reach. Use of commercial cloud-based infrastructure can eliminate much of the local computing cost for this application, and on-demand scalability is expected to allow decreased calculation turnaround time without incurring greater cost. Security concerns are likely to be solvable, and costs are likely to reduce. We conclude that cloud computing is an efficient and cost effective way of simulating radiation transport in a patient, making it attractive as a treatment planning or plan validation tool for any radiotherapy modality.

1. Introduction

Proton beam therapy is a radiotherapy modality whereby accelerated protons are used to deliver dose to tumours while sparing surrounding healthy tissue. Proton therapy has particular advantages over conventional X-ray radiotherapy in cases where the integral dose to healthy tissue must be minimised, for example in paediatric cancers (MacDonald et al. 2011). By the end of 2013 more than 100,000 patients had been treated with proton beam therapy worldwide (PTCOG 2014), and proton beam therapy is rapidly expanding to treat ever more patients each year. While the dose distribution offered by proton therapy is superior, the full benefits cannot yet be realised due to limitations in both patient imaging and treatment planning. Treatment planning for proton therapy makes use of analytic or semi-analytic dose models, in particular the so-called pencil beam models (Petti 1992, Hong et al. 1996, Hollmark et al. 2004, Bortfeld 1997) which are used in the majority of commercial treatment planning codes. Such models are used because they are fast, typically taking less than one minute to calculate the dose distribution depending on the complexity of the plan. However, the analytical models used in such codes have several significant drawbacks, notably in their treatment of multiple Coulomb scattering; this can lead to significant inaccuracies when estimating dose deposition (for example in a strongly inhomogeneous patient geometry, particularly when incident beams traverse parallel to those density inhomogeneities (Paganetti 2012)).

One method of improving the accuracy of the dose deposition estimate is to use Monte Carlo (MC) simulation - in which individual simulated primaries are used to statistically estimate the dose deposited in the patient - to validate a radiotherapy plan or, if fast enough, to produce the radiotherapy plan directly using MC simulation. In recent years, many attempts have been made to produce fast Monte Carlo codes which can produce a dose profile on a timescale rivalling that of the semi-analytical calculations. Recent publications have shown this being done by rewriting the simulations to utilise (for example) Graphical Processing Units (GPUs) (Jia et al. 2012). To achieve the desired calculation speed these codes typically make sacrifices in the accuracy of their models, for example by simplifying either the number of materials modelled, by omitting some physical processes, or by using approximations in the models for scattering and straggling (Kohno et al. 2002, Jia et al. 2014).

To take advantage of a GPU code must often be written in a non-portable way, typically requiring the use of a proprietary and platform-specific language such as CUDA (Nickolls et al. 2008). Platform-specific codes not only need a re-implementation of physics models, but also careful validation of those model implementations against measurements (Jia et al. 2012, Jia et al. 2014).

In this work we present a simple and effective method to obtain full physics validation of a mock pencil beam scanning (PBS) treatment plan using simulations built on a well-validated and portable code base. Simulation times approaching those of the semi-analytical models are achieved through the use of cloud computing to carry out GEANT4 simulations of the treatment plan in under 10 minutes. These simulations were benchmarked against a local many-core cluster and for reference against a single-core desktop machine. Comparisons of the dose distributions produced in the cloud to those from a local cluster show the distributions to be entirely equivalent.

2. Use of GEANT4

GEANT4 is a general-purpose Monte Carlo framework originally developed for use in high-energy physics simulation; however it has recently found application in the medical physics community, particularly for radiotherapy dose calculations. GEANT4 is a well-validated code (Ahn et al. 2013, Lechner et al. 2010, Pablo G A Cirrone & Russo 2006) with tuneable physics implementations that allow the user to strike a balance between the speed and accuracy of simulations. Reducing the number of phenomena considered can increase the calculation speed, but with a reduction in accuracy. Validation in GEANT4 is done in two ways: validation of microscopic quantities, such as cross sections, angular/energy distributions and other parameters; validation of macroscopic quantities such as particle range, straggling and scattering behaviour, which often involve the simulation of an entire experimental setup (Amako et al. 2006). GEANT4 implements chosen sets of physical phenomena using the so-called physics lists (Agostinelli et al. 2003) - these are pre-defined sets of models which are deemed accurate enough for a given simulation need. In proton beam therapy the physics lists needed to achieve acceptable accuracy were explored by Grevillot (Grevillot et al. 2010) and later refined by Kurosu (Kurosu et al. 2014) for the GEANT4-based program GATE (Jan et al. 2004). These settings have been used in the benchmark simulations here to give a realistic timing for a representative mock proton therapy treatment plan.

To carry out a validation of a treatment plan using Monte Carlo - for example to assess the accuracy of a dose distribution obtained from a treatment planning system that uses a semi-analytical algorithm - we must simulate a sufficient number of primary incident protons drawn from a set of beam angles and positions that together define the treatment plan and its beam sources. For the test plan used in this study 10^7 primary protons are needed to obtain an uncertainty of roughly 1% in the estimate of dose deposited within the 95% isodose surface, using the method described in Chetty et al. (2006) to compute the dose uncertainties. A treatment plan validation simulation requires a representation of the patient geometry, e.g. a CT image of the patient to be treated, and a description of the treatment itself (beam energies, angles, spot weightings etc.). Several files are also used to perform administrative functions, for example to ensure statistical independence of the separate Monte Carlo simulations used to determine a single dose estimate. The CT image file may be converted into GEANT4 materials using the stoichiometric calibration presented by Schneider et.al. (Schneider et al. 2000) and the NIST (National Institute of Standards and Technology) database within GEANT4 for individual material properties. The CT image used here was of an anatomical phantom developed at the Christie Hospital - known as MARVIN (Aitkenhead et al. 2013) - digitised as $222 \times 425 \times 171 \approx 16$ million voxels each of $0.976 \times 0.976 \times 2.0$ mm^3 and spanning a total volume of approximately $21.6 \times 41.4 \times 34$ cm^3. Dose distributions were produced as 3D histograms recording the dose deposited in dose grid voxels (dosels), and written to binary format files for later processing.

A mock treatment plan was devised using the RayStation treatment planning suite (RaySearch Laboratories, Stockholm, Sweden) version 2.9, and translated into GEANT4 macro files. One macro file defines the image name and the properties of the required output dose grid, and performs Monte-Carlo-specific initialisations such as random number seeding. A second macro file defines the position, direction, spot size, energy and intensity of each spot in the proton PBS beams to replicate the treatment plan in GEANT4. To mimic a clinical proton treatment plan for a brain tumour, two

Figure 1. A picture of MARVIN, a radiotherapy verification phantom developed at the Christie Hospital. The phantom is derived from a series of patient contours taken from CT images which are used to construct an average head and neck shape. The phantom is constructed in ABS, with removable inserts to provide cavities within the phantom for both ionisation chambers and film.

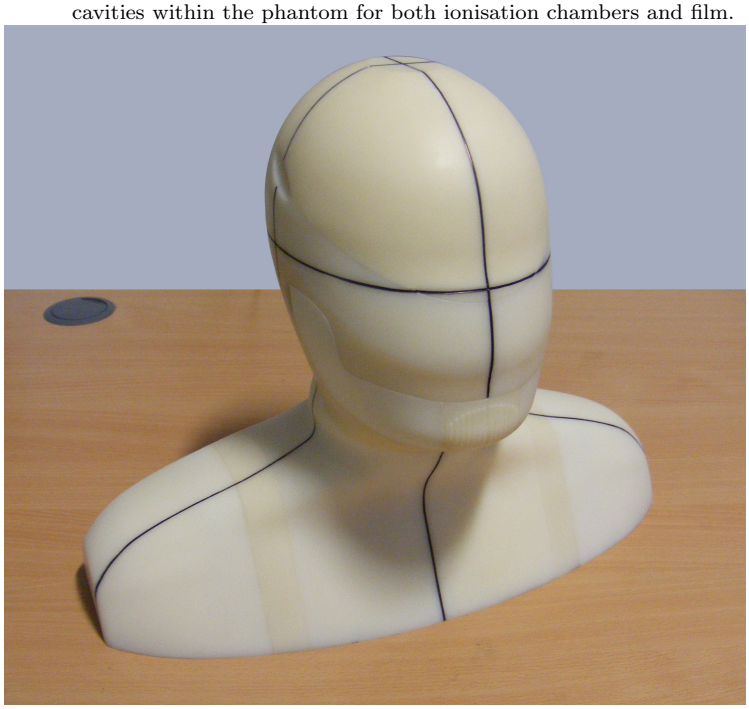

opposing Single Field Uniform Dose (SFUD) fields were used to obtain the desired spherical dose distribution. The fields contain 1514 spots from the right and 1470 spots from the left, and produce a sphere of dose whose centre is around 7cm below the surface of the head; the target volume was roughly 50cm^3. In this plan, proton energies range from 30 to 132 MeV. In total these macros require 1.3MB of disk space, and the CT image requires a further 31MB; a compressed file archive containing the macros and image requires 5MB. Dose distributions were produced by binning the dose deposited in $2 \times 2 \times 2$ mm^3 dosels; this resolution was chosen to be representative of common dose distribution resolutions used in clinical practice. The dose distribution image covered the entire phantom volume, and contained $108 \times 208 \times 171 \approx 4$ million dosels.

3. Computer Architectures

3.1. Local Computing

GEANT4 version 10.00.p02 was used due to the introduction in this version of the multi-threaded event processing needed (Ahn et al. 2013); this means that several incident protons may be tracked in parallel, with the number tracked dependent on the number of processor cores or threads available. The use of multi-threaded GEANT4 allows for reductions to be made in the amount of RAM required to perform a

simulation. In previous versions, multiple processes would be spawned each with their own memory space duplicating simulation data (for example the geometry description) which does not differ from one thread to another. The new multi-threaded version dispenses with this duplication.

To provide a comparative baseline against which to assess the multiprocessor systems, a simulation using a single thread on a consumer desktop computer was performed. This was a Dell workstation with an Intel Core2 Duo dual-core processor running at 3.00 GHz. Such a computer costs less than $1560 (£1000) when new, and has a power consumption of around 500W.

As an example of a typically-available cluster computer, a locally-hosted server was used, comprising 4 AMD Opteron 6174 processors running at 2.2 GHz giving a total of 48 cores. This machine draws roughly 1kW under full load; such a system costs approximately $8400 (£5400).

3.2. The Google Compute Engine

One of the motivations to use cloud computing for treatment plan validation is to remove the need for local computing hardware, and thereby have a reduction in capital cost and infrastructure. All simulations performed in this work were launched and dispatched over a network connection with a nominal 100Mbit/sec upload and download speed. The Google Compute Engine (GCE) service was chosen as the cloud provider for two main reasons - ease of application development, and competitive pricing structure.

The servers provisioning a GCE instance - a virtual machine running on the GCE infrastructure - are Intel-based and use Xeon E5 processors running at 2.6GHz with up to 16 threads per processor. The GCE product offering defines several machine configurations suited to different use cases; the number of threads and RAM available varies with the machine configuration used, as does the price. In this work the n1-highcpu-16 configuration was used, consisting of 16 threads with 14.40GB RAM per machine instance. 100 machine instances were created, giving a total of 1600 cores; each machine was then used to calculate the dose from $1/100^{\text{th}}$ of the total number of required primary protons. The resulting dose distribution was then remotely combined into a single distribution on one of the cluster instances. GCE usage is billed in per-minute increments, and varies for the different machine configurations available. The n1-highcpu-16 configuration used here costs $0.668 (£0.43) per hour per machine instance with a minimum charge of 10 minutes' usage, although this price is predicted to fall sharply.

To minimise the initialisation time for a validation simulation, a customised disk image was created and stored in the cloud. This disk image is copied into virtual hard drives when a server is launched and used to boot each system. The image contained an operating system, an installation of GEANT4, a copy of the GEANT4 application used to perform the simulation and utilities to facilitate combination of dose distributions in the cloud. Cloud storage of this image costs $0.085 (£0.05) per GB per month - an insignificant amount. Cloud storage is used to distribute the treatment plan and associated CT images to the servers, and also to hold the returned dose profiles before downloading; this is also hosted by Google and currently costs $0.026 (£0.02) per GB per month. Network traffic to and from the Google cloud is presently charged at a rate of $0.12 (£0.07) per GB. Fast deployment was made possible through the use of the GCE replica pools API (Google 2014*b*), which allows

users to create a pool of servers with identical properties. Random number seeds were derived from the unique server instance IDs to ensure statistical independence of the simulations, and these seeds were transferred to the cloud storage to facilitate repeat running if needed. Timing was achieved by date-stamping the start and end of the node lifetime, and transferring this data to the cloud storage for later analysis. The cluster is instantiated using a Python script on the local client to access the GCE API through JSON http commands; this allows very fast deployment and repeatability.

For all but the desktop computer benchmark, data has to be transferred over a network. For the local cluster this was over a standard local 100Mbit/sec Ethernet connection and achieved transfer speeds of 14MB/s in either direction. Transfers to the cloud computers have to take place over the public Internet, and passed through the institutional broadband connection of the University of Manchester. This connection allowed transfers to and from the cloud service at approximately 3MB/sec, though this varied considerably ranging from 2MB/sec up to 11MB/sec. After the initial upload of a plan data can be transferred entirely within the cloud where the transfer speeds are much higher. Given the measurements presented below, we do not believe network speed presents a significant bottleneck; indeed, it was possible to launch benchmarks using a domestic broadband connection without significant impact on computation times.

4. Results

Figure 2 indicates the dose distribution produced by the mock plan. In this figure, the dose distribution is shown overlaid on a CT image of the MARVIN phantom. As described above, the dose distribution is spherical with a radius of 5 cm, and lies at a depth of approximately 7 cm inside the head of the phantom.

Figure 2. Sample grids overlaid on a CT image of the MARVIN phantom. The size and shape of the high-dose region can be seen, along with some of the internal features of the MARVIN phantom. Deviation from a spherical dose is caused by small inhomogeneities in the phantom which are then seen in the Monte Carlo simulation. Each dosel is $2 \times 2 \times 2$ mm^3.

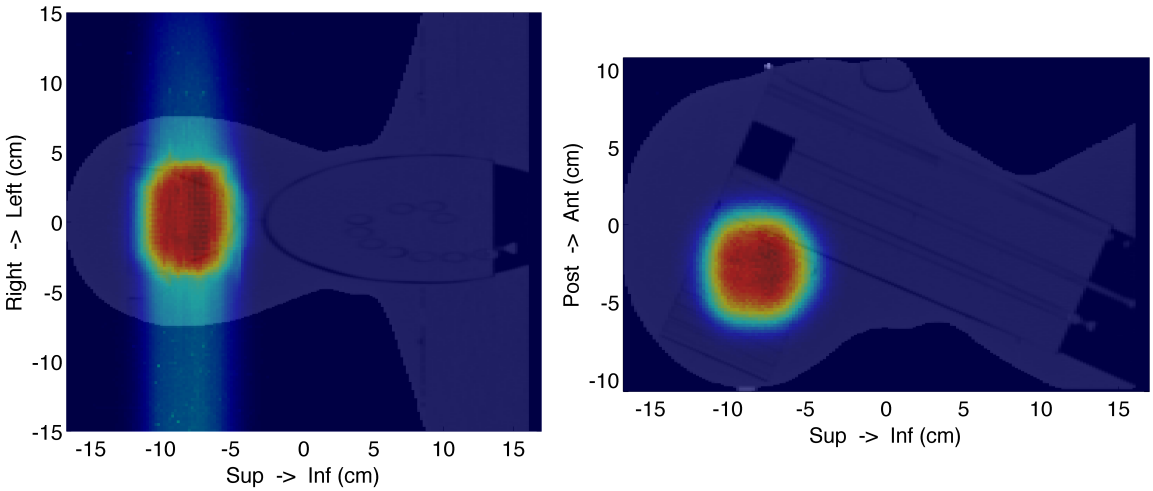

To assess the reliability of the dose distributions computed in the cloud, a Gamma analysis (Low et al. 1998, Low & Dempsey 2003) was performed between the cloud dose distribution and the dose distribution produced by the local cluster. The result of this analysis when a Gamma criterion of 3%/3mm is applied is shown in Figure 3. At the Christie Hospital the criteria for clinical Gamma analyses are dependent on treatment site, but 3%/3mm is typical with the requirement that 95% of all points having a dose greater than 20% of the maximum dose must have a Gamma index less than 1. The pass rate of the dose distributions shown in Figure 3 is 97.64%, meaning it passes the clinical criteria. Further Gamma analysis pass rates are given in Table 1, including comparisons with stricter criteria. The sharp fall off in Gamma analysis pass rate when the criteria is tightened to 1%/1mm is due to the increasing dominance of statistical fluctuations in the dose distributions, and could be rectified by running more particle histories to reduce the statistical noise.

Figure 3. Gamma analysis between the cloud-produced and the locally-produced dose distributions. Plots A and B show the dose distributions in the regions with dose higher than 20% of the peak dose, while the rightmost plot shows the Gamma index map. Points whose dose is too low are coloured blue - points whose dose is too high are red; passing points are green. A global dose-difference method was applied in the Gamma analysis.

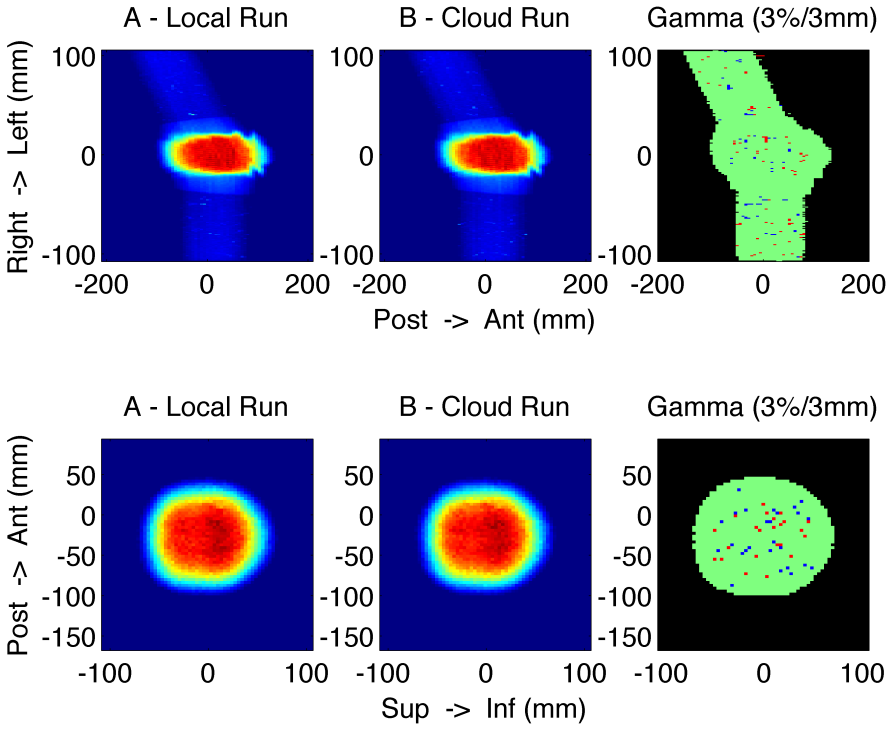

Table 2 shows the time in seconds for each simulation to run. Unsurprisingly, the single-threaded simulation on the desktop machine took the longest time (31h32m43s), and serves to illustrate the approximate number of CPU seconds required for one simulation. The local cluster completed the simulation in approximately 80 minutes

Table 1. Gamma analysis pass rates for various criteria levels. The pass rate trend is as expected - tighter criteria result in lower pass rates. The sharp drop between the 2%/2mm and 1%/1mm results is due to the increasing influence of statistical noise, which is about 1% for this case.

Analysis Criteria	Gamma pass rate
4%/4mm	98.90%
3%/3mm	97.64%
2%/2mm	93.26%
1%/1mm	74.17%

(1h20m46s). The clear winner in terms of raw simulation speed is the cloud simulation (3m21s) but it should be remembered that this does not include the time taken to combine and download the output data, which is an extra step not required with the other machines.

Table 2. Times taken to simulate 10 million protons on the three tested computer systems. These times reflect only the simulation time, and not the time taken to combine or download the data where performed.

Computer	Time (s)	Time (h/m)	Comments
Desktop - 3.0 GHz Intel Core2 Duo	113563	31h 32m 43s	Single thread
Cluster - 2.2 GHz AMD Opteron 6174	4846	1h 20m 46s	48 threads
Cloud cluster - 2.6 GHz Intel Xeon E5	201	3m 21s	16 threads per node, 100 nodes. Maximum time, not including combination of data or download

In Table 3, the transfer and combination times for the cloud simulation are shown, with approximate transfer speeds indicated. All data is stored in the Google cloud and transfers are performed using the Google cloud API tool 'gsutil' (Google 2014*d*). This tool allows parallel transfers to and from the cloud storage, which vastly improves the download speed. The upload of the plan data described in Section 2 is only done once, and it is distributed to the machines on startup.

Post-processing time was measured independently for simulation data produced in the cloud and the average, maximum and minimum times for 10 post-processing runs were recorded. The average total time spent in post-processing data from the cloud cluster to produce a dose grid equivalent to that from the local cluster machine is 47.91 +12.07 -13.18 seconds (where +12.07 and -13.18 indicate maximum and minimum times respectively), meaning that the minimum total time to complete the simulation, including all data transfer and combination is 246.69 + 12.54 - 6.91 seconds - i.e. 4m7s on average. The actual time taken depends on factors including the node launch time, statistical variations in the node run time and the particular method used to trigger the combination of data. The total time taken for this example calculation was 6m34s from start to finish, including all launch times, transfers and post processing.

Table 4 shows the approximate cost of producing a validation using each of the machines in this work, without the amortised cost of hardware included. The local

Table 3. Network transfer and post-processing times for the simulations performed in the GCE, showing the mean value together with the maximum and minimum times derived from 10 separate post-processing runs over the same simulation data. Dose grids were transferred to a 'master' node, where they were combined into a single image to be compressed and uploaded to the cloud storage. The subsequent download of this combined image was also timed.

Step	Time taken (s)	Transfer speed (MB/s)	Notes
Upload of plan data	$3.45^{+1.78}_{-0.99}$	$0.70^{+0.4}_{-0.1}$	Upload size: 5MB
Collation of dose data	$42.37^{+10.55}_{-12.38}$	$68.44^{+28.26}_{-13.64}$	Transfer 2.9GB from cloud storage to master instance
Combination of dose data	$1.74^{+0.14}_{-0.12}$	N/A	Final dose image size: 29MB
Download of dose data	$3.80^{+1.38}_{-0.68}$	4.91±0.3	Compressed, total download size: 18.65MB
Total post processing	$47.91^{+12.07}_{-13.18}$		

cluster is the cheapest option, but takes much longer than the cloud-based validation. The price quoted for the GCE is that of 100 n1-highcpu-16 machines for the minimum billing time of 10 minutes - an artificial billing constraint that in principle could be removed. If this minimum time could be abolished, the price would drop to $4.65 (£2.98) which would make the GCE validation very competitive; in any case, prices for the GCE instances fell 13% during the course of this work (Google 2014*a*), and are expected to fall further in the coming months.

Table 4. Cost per validation using the three machines tested in this work. The cost of electricity has been taken to be 15p/kWhr and power consumptions from the earlier specifications are used. It is important to note that other costs, such as hardware renewal and cooling are not included in these figures. The cost of the GCE instances is based on the current 10-minute minimum charge levied by Google - this cost could fall to $4.65 (£2.98).

Computer	Time taken	Hardware/Hire cost	Total running cost (per validation)
Desktop	31h 32m 43s	$1560 (£1000)	$3.70 (£2.37)
Local cluster	1h 20m 46s	$8400 (£5400)	$0.31 (£0.20)
GCE	4m 7s	$0.688/hour (£0.44/hour)	11.20^{*} (£7.17)*

5. Discussion

MC calculation of treatment plans was significantly faster in the GCE than in the local cluster (by a factor of 24) and the single desktop computer (by a factor of 565) used in this work, but is slower than fast MC codes such a gPMC (Jia et al. 2012). However, cloud-based usage of GEANT4 has a number of advantages over GPU codes: the relative ease of validation and modification of physics models, better out-of-field dose accuracy due to superior physics modelling; and transparency of porting the code to the cloud. The use of 'vanilla' GEANT4 means that other codes designed to simulate radiotherapy applications - for example GATE (Jan et al. 2004) - may also be used in the cloud, whereas codes like gPMC have specialised applications and require specific hardware.

A high speed network is required to efficiently transfer data to and from the cloud. It should be noted that within the cloud, all dose data transfers in this study were done without compression meaning around 3GB were transferred inside the cloud during each validation. Enabling compression would reduce the amount of data to transfer by a factor of two, possibly even more. However, within-the-cloud compression would likely result in a lower total throughput due to the time required to compress and decompress the data. When downloading the data to a local machine compression is beneficial and results in a transfer speed improvement of around 1.4× over an uncompressed transfer. Given the transfer times stated above, and the general tendency for network speeds to increase, we believe the network is not a significant bottleneck.

Data transfer and its security may be an issue, specifically the protection of patient confidentiality. Guidance provided by a UK government report into patient data protection is unclear on this, with most of its recommendations being concerned with sharing patient information for scientific studies and other similar purposes (Caldicott et al. 2013). It is our belief that the encryption used in the transfer and temporary storage of data in the cloud is sufficient to meet required security standards in at least the UK, and if needed could be complemented by additional encryption or anonymisation techniques (Google 2014*c*). Given the good speed of the validation shown here, this is a topic which deserves investigation.

6. Conclusions

The test simulation completed 24 times faster in the GCE than in the local cluster and 565 times faster than in a single desktop computer with no compromise in accuracy for a cost of $11.20 (£7.17). Computation of dose grids in the GCE is a promising alternative to the use of a local cluster (whether conventional or GPU architecture), with job turnaround speed being a significant potential benefit. While the time taken to produce the dose grid in the cloud is impressive, this technique cannot readily rival the speed of GPU-based codes such as gPMC and others for the purpose of in-field dose calculation; however the improved accuracy of the physics models provided by the GEANT4 toolkit should be an incentive for the use of this method. It is expected that the price per CPU hour in cloud computing will fall faster than for locally-based computing (Google 2014*a*), and provided the presently (rather arbitrary) 10 minute minimum charge can be removed, the cloud will become ever more competitive in terms of price per validation as well as in terms of speed. We believe that issues around the network speed required are not significant, with the potential use of compression and

the ever-increasing speed of internet connections mitigating against the requirement to transfer large amounts of data. The question of the legality of dose computation in the cloud should be looked into, as the potential benefit is large. It is our opinion that the encryption used in the transfer and storage of data is sufficient to comply with legal requirements, for example that data is kept within the European Union (*The Data Protection Act* 1998).

Acknowledgements

Work supported in part by the UK Science and Technology Facilities Council, STFC grant number ST/G004277/1, and by Christie Medical Physics and Engineering.

References

Agostinelli S, Allison J, Amako K, Apostolakis J, Araujo H, Arce P, Asai M, Axen D, Banerjee S, Barrand G, Behner F, Bellagamba L, Boudreau J, Broglia L, Brunengo A, Burkhardt H, Chauvie S, Chuma J, Chytracek R, Cooperman G, Cosmo G, Degtyarenko P, Dell'Acqua A, Depaola G, Dietrich D, Enami R, Feliciello A, Ferguson C, Fesefeldt H, Folger G, Foppiano F, Forti A, Garelli S, Giani S, Giannitrapani R, Gibin D, Gómez Cadenas J J, González I, Gracia Abril G, Greeniaus G, Greiner W, Grichine V, Grossheim A, Guatelli S, Gumplinger P, Hamatsu R, Hashimoto K, Hasui H, Heikkinen A, Howard A, Ivanchenko V, Johnson A, Jones F W, Kallenbach J, Kanaya N, Kawabata M, Kawabata Y, Kawaguti M, Kelner S, Kent P, Kimura A, Kodama T, Kokoulin R, Kossov M, Kurashige H, Lamanna E, Lampén T, Lara V, Lefebure V, Lei F, Liendl M, Lockman W, Longo F, Magni S, Maire M, Medernach E, Minamimoto K, Mora de Freitas P, Morita Y, Murakami K, Nagamatu M, Nartallo R, Nieminen P, Nishimura T, Ohtsubo K, Okamura M, O'Neale S, Oohata Y, Paech K, Perl J, Pfeiffer A, Pia M G, Ranjard F, Rybin A, Sadilov S, Di Salvo E, Santin G, Sasaki T, Savvas N, Sawada Y, Scherer S, Sei S, Sirotenko V, Smith D, Starkov N, Stoecker H, Sulkimo J, Takahata M, Tanaka S, Tcherniaev E, Safai Tehrani E, Tropeano M, Truscott P, Uno H, Urban L, Urban P, Verderi M, Walkden A, Wander W, Weber H, Wellisch J P, Wenaus T, Williams D C, Wright D, Yamada T, Yoshida H & Zschiesche D 2003 *Nuclear instruments & methods in physics research. Section A, Accelerators, spectrometers, detectors and associated equipment* **506**(3), 250–303.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0168900203013688*

Ahn S, Apostolakis J, Asai M, Brandt D, Cooperman G, Cosmo G, Dotti A, Dong X, Yung Jun S & Nowak A 2013 *in* 'SNA + MC 2013' pp. 1–8.

Aitkenhead A H, Rowbottom C G & Mackay R I 2013 *Physics in medicine and biology* **58**(19), 6915–6929.

Amako K, Guatelli S, Ivanchencko V, Maire M, Mascialino B, Murakami K, Pandola L, Parlati S, Pia M G, Piergentili M, Sasaki T & Urban L 2006 *Nuclear Physics B - Proceedings Supplements* **150**, 44–49.

Bortfeld T 1997 *Medical physics* **24**, 2024–2033.

Caldicott F, Carvel J, Catchpole M, Dafter T, Davies J, Haslam D, Hassey A, Monaghan D, Parkin T, Partridge N, Severs M, Tapster C, Taylor J, Walport M & Wrigley D 2013 The Information Governance Review Technical report.

Chetty I J, Rosu M, L K M, Fraass B A, Ten Haken R K, Kong F M & McShan D 2006 *Radiation Oncology Biol. Phys* **65**(4), 1249–1259.

Google 2014*a* 'Announcing across-the-board price cuts on Compute Engine'.
URL: *http://googlecloudplatform.blogspot.co.uk/2014/10/Announcing-across-the-board-price-cuts-on-Compute-Engine.html*

Google 2014*b* 'Google Cloud Platform Replica Pools API'.
URL: *https://cloud.google.com/compute/docs/replica-pool/v1beta1/*

Google 2014*c* 'Google Security Audits and Certifications'.
URL: *https://services.google.com/fh/files/blogs/btd-sec-op-2014-grey.pdf*

Google 2014*d* 'gsutil tool'.
URL: *https://cloud.google.com/storage/docs/gsutil*

Grevillot L, Frisson T, Zahra N, Bertrand D, Stichelbaut F, Freud N & Sarrut D 2010 *Nuclear*

Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **268**(20), 3295–3305.

Hollmark M, Uhrdin J, Belki D, Gudowska I & Brahme A 2004 *Physics in medicine and biology* **49**(14), 3247–3265.

Hong L, Goitein M, Bucciolini M, Comiskey R, Gottschalk B, Rosenthal S, Serago C & Urie M 1996 *Physics in medicine and biology* (41), 1305–1330.

Jan S, Santin G, Strul D, Staelens S, Assie K, Autret D, Avner S, Barbier R, Bardies M, Bloomfield P, Brasse D, Breton V, Bruyndonckx P, Buvat I, Chatziioannou A F, Choi Y, Chung Y H, Comtat C, Donnarieix D, Ferrer L, Glick S J, Groiselle C J, Guez D, Honore P F, Kerhoas-Cavata S, Kirov A S, Kohli V, Koole M, Krieguer M, van der Laan D J, Lamare F, Largeron G, Lartizien C, Lazaro D, Maas M C, Maigne L, Mayet F, Melot F, Merheb C, Pennacchio E, Perez J, Pietrzyk U, Rannou F R, Rey M, Schaart D R, Schmidtlein C R, Simon L, Song T Y, Vieira J M, Visvikis D, Van de Walle R, Wieers E & Morel C 2004 *Physics in medicine and biology* pp. 4543–4561.

Jia X, Schümann J, Paganetti H & Jiang S B 2012 *Physics in medicine and biology* **57**(23), 7783–7797.

Jia X, Ziegenhein P & Jiang S B 2014 *Physics in medicine and biology* **59**, R151–R182.

Kohno R, Sakae T, Takada Y, Matsumoto K, Matsuda H, Nohtomi A, Terunuma T & Tsunashima Y 2002 *Japanese Journal of Applied Physics* **41**(Part 2, No. 3A), L294–L297.
URL: *http://stacks.iop.org/1347-4065/41/L294*

Kurosu K, Takashina M, Koizumi M, Das I J & Moskvin V P 2014 *Nuclear Inst. and Methods in Physics Research, B* **336**(C), 45–54.

Lechner A, Ivanchenko V N & Knobloch J 2010 *Nuclear Inst. and Methods in Physics Research, B* **268**(14), 2343–2354.

Low D A & Dempsey J F 2003 *Medical physics* **30**(9), 2455.

Low D A, Harms W B, Mutic S & Purdy J S 1998 *Medical physics* **5**(25), 656–661.

MacDonald S M, Trofimov A, Safai S, Adams J, Fullerton B, Ebb D, Tarbell N J & Yock T I 2011 *Radiation Oncology Biology* **79**(1), 121–129.

Nickolls J, Buck I, Garland M & Skadron K 2008 *Queue* **6**(2), 40–53.
URL: *http://doi.acm.org/10.1145/1365490.1365500*

Pablo G A Cirrone G C F D R S G B M M G P & Russo G 2006 *in* 'IEEE Nuclear Science Symposium Conference Record' pp. 788–792.

Paganetti H 2012 *Physics in medicine and biology* **57**(11), R99–R117.

Petti P L 1992 *Medical physics* **19**(1), 137.

PTCOG 2014 Patient Statistics - March 2014 Update Technical report PTCOG.

Schneider W, Bortfeld T & Schlegel W 2000 *Physics in medicine and biology* (45), 459–478.

The Data Protection Act 1998 *The Stationery Office Limited, UK* p. c.26.