Measuring radiation treatment plan similarity in the cloud

J. Andrea, C. Pinter and G. Fichtinger

Laboratory for Percutaneous Surgery, School of Computing, Queen's University, Kingston, Canada

Abstract— **Radiation therapy is a form of cancer treatment in which carefully designed plans are used to direct treatment over multiple occasions (fractions). Creating radiation plans is quite laborious, so it is not feasible to manually create a plan for each fraction to maintain treatment quality. We propose to use a database of plans to find the most similar anatomy, based on which a suitable daily plan might be automatically created, thus reducing staff time. However, the computation for finding the most similar plan is long and computationally intensive. We present a method for finding the most similar plan using cloud resources to reduce computation time. The preliminary, unoptimized implementation finds the most similar plan in less than 7 minutes when choosing from five plans, and it is anticipated that increasing the number of plans will result in only a small relative increase in computation time.**

I. Introduction

Radiation therapy is a form of cancer treatment in which carefully designed plans are used to direct treatment. The radiation plan is computed prior to treatment in order to administer the maximum dose possible to the target and the minimum dose possible to any nearby organs at risk. At the beginning of the treatment planning process, the patient is scanned, typically with computed tomography (CT). On the resulting CT scan, contour lines are manually or semi-automatically drawn around all structures of interest. The optimized radiation treatment plan can be created using the CT scan, the contours, and the prescription.

The radiation dose must be given in fractions over a period of several days because giving the total radiation dose in one treatment causes too much damage to nearby healthy tissue. When creating a new treatment plan, optimizing a plan previously created for this patient rather than a generic plan may reduce the optimization time. However, the patient's anatomy frequently differs from fraction to fraction due to anatomical changes such as weight loss or an empty stomach [1] [2]. Choosing the treatment plan previously computed for this patient with the most similar anatomy to the patient's daily anatomy as an "initial guess" for further optimization would be a more appropriate choice [3]. A similar approach was presented by Kazhdan *et al.* in 2009 [3].

However, the procedure for determining the most similar plan is computationally intensive, and can take a long time depending on how many plans have been created for this patient, which presents an obstacle to performing the procedure clinically. As there is limited time between acquiring the CT scan and starting treatment, the most similar plan must be found quickly, as it requires further optimization and must be checked and approved by a physician.

In cloud computing, resources are shared by users in order to maximize their effectiveness, and are reallocated based on demand. Using the cloud avoids upfront infrastructure costs, as users only pay for the cloud resources they use. The cloud provides elasticity and scalability of resources through dynamic allocation, and is also low-maintenance. Many different applications currently take advantage of cloud resources, including e-commerce, web applications like Facebook and LinkedIn [4], crowdsourcing, and offloading code from mobile devices [5]. These applications benefit from usage of the cloud, including reducing battery usage and storage space, and elastically scaling resources as application usage fluctuates.

For this work, two techniques were designed and implemented. The first is a workflow using open source software for selecting the most appropriate radiation treatment plan to be optimized which was previously computed for the patient, based on similarity between the structures of interest when the initial plan was computed and the same structures for the current fraction. The second was to perform this computation in the cloud. The two techniques in combination comprise a component of the full system needed for treatment plan selection and registration [6]. The two components implemented in this paper are referred to as the treatment plan selection system.

II. Methodology

A. Hardware, software, and data

The local computer in this research work is a desktop computer with an Intel Core 2 Quad CPU and 8GB of RAM, running 64-bit Windows 7. It had a clock speed of 2.40 GHz. The local computer was running on a LAN network which had a download speed of 94.89 Mbps and an upload speed of 20.86 Mbps. The local computer was used as the client computer in the treatment plan selection system on the cloud. The cloud instances were Elastic Cloud Compute (EC2) c3.xlarge instances, which are compute-optimized instances, and were running Windows Server R2

D.A. Jaffray (ed.), *World Congress on Medical Physics and Biomedical Engineering, June 7-12, 2015, Toronto, Canada*,
IFMBE Proceedings 51, DOI: 10.1007/978-3-319-19387-8_105

2008 Datacenter. The instances had 4 virtual CPUs, 8GB of RAM, and had Intel Xeon E5-2680 v2 processors. The instances had clock speeds of 2.80 GHz.

Both the client computer and the cloud instances used 3D Slicer and SlicerRT. 3D Slicer is an open source platform for medical image analysis and visualization [7]. SlicerRT is an extension for 3D Slicer that provides tools for radiation therapy research [8]. The instances used this platform to perform the analysis steps, while the client computer used it for visualization of the best plan and any further analysis. Boto [9], which is a python-based interface for Amazon Web Services (AWS), was used to provide coordination between the client computer, the cloud instances, and the specific AWS products used. The AWS products used were the core components for providing the ability to perform the plan selection on the cloud. EC2 instances were used to perform the computation on the cloud, Simple Storage Service (S3) was used to store the radiation treatment plans, the daily CT and structure set, and the results. Simple Queue Service (SQS) was used to provide communication between the cloud instances and the client computer.

The data that was used in this project was simulation data generated for this purpose. It was generated by applying random deformation fields to a radiation plan for a phantom dataset provided by the Cancer Centre of Southeastern Ontario at the Kingston General Hospital. The random deformation fields used displacement vectors in the range -10 mm to 10 mm in all directions. We used a coarse field that assigns the vectors every 125 mm in the anterior-posterior and left-right directions, and every 57.5 mm in the inferior-superior direction.

B. Architecture

In the architecture of the treatment plan selection system, the storage, computation (performed on the instances), and messaging are all done on the cloud. These actions performed by the local computer and the cloud instances include putting data in the cloud storage, including the instances' results and the daily CT and structure set, accessing data from the cloud storage, and messaging between the local computer and the instances via the message queue.

While only a single queue was used in this system, each message put in the queue had an instance ID attached to it to ensure that each message was correctly received by its intended target. When the instances or the local computer were waiting for a message, when they received a message they first checked the instance ID, and released the message back to the queue if it did not match their own ID.

C. Workflow

At the beginning of the workflow, a new treatment plan is created and uploaded, along with the CT scan and contours that were used to create the plan, into the cloud storage. At a later date, the patient comes in for the daily radiation treatment plan fraction. For this proof-of-concept work, the patient's daily cone-beam CT (CBCT) is contoured to delineate the structures of interest.

The technician initiates the treatment plan comparison system, and the client computer connects to the cloud storage and the cloud message queue. It finds the daily study, packages it up and uploads it to the cloud storage. It then sends a message to each instance via the queue notifying them that the daily study has been uploaded, and also sends them each a message informing them which radiation study they will use for the comparison.

When the instances receive the messages with their instance UID, they download their assigned study and the daily study from the cloud storage. They perform their component of the treatment plan selection process, which is described in greater detail later in the paper. Each instance performs a single comparison between their assigned study and the daily study, which produces a similarity value as the result. When the instances finish the comparisons, the results are uploaded to the cloud storage, and the instances each send a message to the client computer that they have finished their computation.

Once the client computer has received a message from each instance that they have finished, it downloads the results from each instance from the cloud storage. The client computer iterates through the results, and chooses the study with the best similarity value. SlicerRT is then launched on the local computer, and the best plan is automatically loaded into SlicerRT for the technician to verify before optimizing it to create the new treatment plan.

D. Comparison computation on the instances

The process of selecting the best treatment plan was automated for this work. Previously, this process could be conducted by manually performing the comparison using the SlicerRT graphical user interface. Each instance performs one comparison between the daily CBCT and structure set and the CT and structure set stored in their assigned study. In the comparison, the daily CBCT scan is registered to the assigned study's CT scan, which produces a rigid transformation matrix that is then applied to the daily structure set. Each pair of matching structures, one from the daily study, and one from the assigned study, is compared. The similarity of the pair of structures is evaluated using the Dice coefficient, which is a statistic for measuring the similarity of two structures. Once all of the structure pairs have been compared, the similarity of the structure sets is computed as the average of the Dice coefficients. This value is output as the similarity value for the assigned study.

As contouring is the most time-consuming step and only happens when a completely new plan is created for the patient, comparison of structures that have been contoured on the CBCT scan was used as a proof-of-concept for using the cloud and SlicerRT together to find the most similar plan. In the future, the similarity measure of raw anatomical data will be used for comparison instead.

E. System evaluation

To verify the accuracy of the system on simulated data, each of the studies stored in the cloud was individually presented to the treatment plan comparison system as the daily study. This provided a ground truth, as the correct results were known. As the limitation with selecting the most similar plan rather than the most recent one is time, it is an important metric for the treatment plan selection process. If the time for the entire selection process is too long, the computation time will present a significant obstacle to adopting this process clinically. For each test, the same five studies in the cloud were used in the treatment plan selection process. Using one study as the daily study, the length of time for performing the comparison process was measured for comparison with one to five studies stored in the cloud, using three different setup configurations. The length of time for performing one comparison between the daily study and an assigned study was also measured for three different setup configurations.

III. RESULTS

The first test performed was to test the treatment plan selection process with each of the five different studies stored in the cloud, where each study was presented to the treatment plan selection system as the daily study. This was a ground truth test for whether the system returned the correct result, as the correct result in this test was the same study stored in the cloud. The system had 100% accuracy, as it returned the correct result every time.

The second test was to measure the length of time required for selecting the treatment plan with the best similarity value using three different setup configurations, with varying numbers of comparison studies. The setup configurations were the new treatment plan selection system on the cloud, which uses an automated comparison process, using the same automated comparison on the local computer, and performing the comparison manually on the local computer. From the results (Fig. 1) it can be seen that while the local

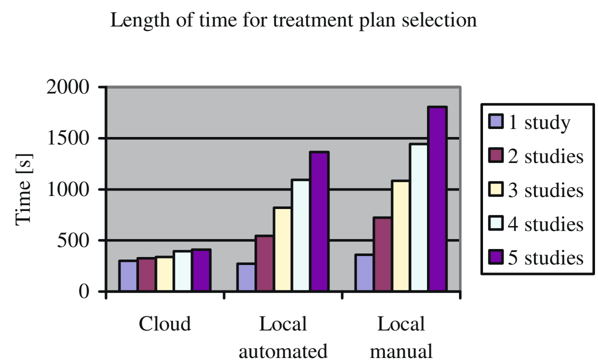

Fig. 1 Length of time for treatment plan selection on different setup configurations. Local automated is the local computer using the automated workflow, and local manual is the local computer using the manual workflow.

computer using the automated comparison performs the fastest for one study, the cloud performs the fastest for two to five studies. For each configuration, the time required increases by approximately the same length of time for that configuration as the number of comparison studies grows.

The next test performed was to measure the time for performing the comparison process once on each of the different setup configurations tested previously. While performing the comparison is the major time component in the treatment selection process on the local computer using either the automated or the manual workflow, the cloud has additional time requirements including uploading data to the cloud and messaging with the client computer. From the results of this test (Fig. 2), it can be seen that the comparison process takes the smallest amount of time when perfor-

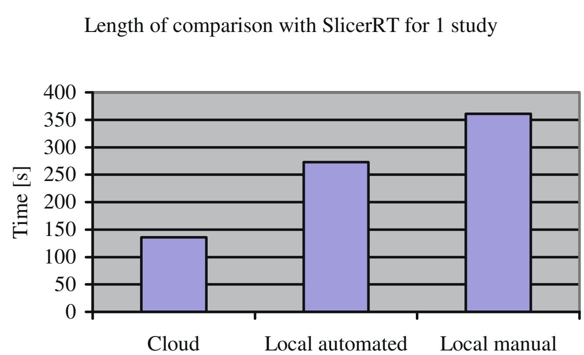

Fig. 2 Length of time for performing one comparison between the daily study and a study containing a treatment plan, using three different configurations. Local automated is the local computer using the automated workflow, and local manual is the local computer using the manual workflow.

med on the cloud, and the largest amount of time when performed on the manual computer.

IV. DISCUSSION

Provided that the increase in time remains consistent as the number of comparison studies grows, we can extrapolate that the cloud will continue to perform the best out of the three for larger numbers of studies (Fig. 1). The increase was approximately 27.5 seconds on average for each new comparison study added, which would suggest that using the treatment plan selection system with ten comparison studies would complete in less than 10 minutes. As treatment plans are not computed frequently for a given patient, the estimated run time is acceptable for this use case.

Performing the comparison between a single study and the daily study took less time on the cloud than either of the other two configurations (Fig. 2). Performing the comparison on the local computer using the manual workflow takes the longest amount of time, so the time increase is due to a human user being slower at performing the workflow than the computer. The smaller amount of time required to perform the comparison on the cloud than to perform it on the local computer using the automated workflow can be attributed to the larger CPU speed of the instances on the cloud. The smaller amount of time also indicates that there is some overhead when using the cloud, as it takes longer to do the treatment plan selection for one study (Fig. 1).

There are some significant advantages with performing the treatment plan selection system on the cloud. The primary advantage is that on the cloud, the comparison between an assigned study and the daily study can be done in parallel. As this comparison time is the largest component of the system in terms of the length of time to completion, being able to perform the comparisons in parallel rather than sequentially on the local computer presents a significant advantage as it greatly decreases computation time.

Due to the system design, adding new studies and therefore new cloud instances to the system requires minimal user effort. This system, while currently set up for five studies, can easily be extended to hundreds or thousands of studies with little effort if other patients' plans are also included. This system has an additional benefit in that the radiation studies are stored on the cloud, reducing local hard drive space, and can be accessed from any location.

V. CONCLUSIONS

In conclusion, this system presents a new use of the cloud with a system for selecting the previously created radiation treatment plan with the greatest anatomical similarity to the patient that day. The selected plan would then be further optimized to create the new treatment plan. The decrease in time required to perform this computation when using the cloud decreases the time to the point where it would be reasonable to perform this computation clinically. This will lead to a reduction in the time required to create a new optimized radiation plan for the same patient.

ACKNOWLEDGMENT

This work was funded by Cancer Care Ontario through Applied Cancer Research Unit and Research Chair in Cancer Imaging and OCAIRO grants.

CONFLICT OF INTEREST

The authors declare that they have no conflict of interest.

REFERENCES

1. Bhide S A, Davies M, Burke K et al. (2010) Weekly volume and dosimetric changes during chemoradiotherapy with intensity-modulated radiation therapy for head and neck cancer: a prospective observational study. International Journal of Radiation Oncology* Biology* Physics 76(5): 1360-1368.
2. Keall P. (2004) 4-dimensional computed tomography imaging and treatment planning. In Seminars in radiation oncology 9(14):81-90.
3. Kazhdan M, Simari P, McNutt T et al. (2009) A shape relationship descriptor for radiation therapy planning. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009: 100-108.
4. Petcu D, Macariu G, Panica S et al. (2013) Portable cloud applications—from theory to practice. Future Generation Computer Systems 29(6): 1417-1430.
5. Elgazzar K, Martin P, Hassanein H. (2014) Cloud-Assisted Computation Offloading to Support Mobile Services. IEEE Transactions on Cloud Computing 99:1.
6. Miller M S, Elizabeth P, Tracton B S E, et al. (1998) Image registration: an essential part of radiation therapy treatment planning. International Journal of Radiation Oncology* Biology* Physics, 40(1), 197-205.
7. Fedorov A, Beichel R, Kalpathy-Cramer J et al. (2012) 3D Slicer as an image computing platform for the Quantitative Imaging Network. Magnetic resonance imaging 30, no. 9: 1323-1341.
8. Pinter C, Lasso A, Wang A et al. (2012) SlicerRT: Radiation therapy research toolkit for 3D Slicer. Medical physics 39(10): 6332-6338.
9. https://github.com/boto/boto

Author: Gabor Fichtinger
Institute: Queen's University
Street: 25 Union St.
City: Kingston
Country: Canada
Email: gabor@cs.queensu.ca