FAMPO
FEDERATION OF AFRICAN MEDICAL PHYSICS ORGANIZATIONS

African Journal of Medical Physics

2022; 4(1): 1-9. doi: 10.

Research Paper

High-throughput, high-availability automated planning for radiotherapy clinics in low-resource settings

Marquez Barbara [1,✉], Gay Skylar [2], Douglas Raphael [2], Zhang Lifei [2], El Basha Daniel [1], Huang Kai [1], Cardenas Carlos E. [3] and Court Laurence E.[1, 2]

1. The University of Texas MD Anderson Cancer Center Graduate School of Biomedical Sciences, Houston, TX, USA.
2. The University of Texas MD Anderson Cancer Center, Department of Radiation Physics, Houston, TX, USA.
3. The University of Alabama at Birmingham, Department of Radiation Oncology, Birmingham, AL, USA.

✉ Corresponding author: Barbara Marquez, MD Anderson Cancer Center, Radiation Physics – Research, 1400 Pressler Street, FCT8.6014, Houston, TX 77030-4009. Tel: +1 (786) 351-2115; E-mail: bmarquez@mdanderson.org

ISSN 2643-5977

Received: 2021.12.05; Accepted: 2022.02.21; Published: 2022.04.01

Abstract

The Radiation Planning Assistant (RPA), a web-based tool for automating the treatment planning process in radiation oncology, can significantly reduce the staff and time needed for treatment planning in low- and middle-income countries. To enable clinical use, the RPA must output radiation plans consistently and reliably, especially as it is deployed to clinics around the world. To test reliability, we performed a thorough capacity study of the RPA for treatment of two disease sites: cervix and head & neck. The RPA architecture consists of several, multi-capacity computing modules (contouring, plan creation, optimization, quality assurance) that process patients serially. Completion times for each module were measured by processing 25 cervix and 25 head and neck (H & N) patient datasets through the entire workflow. Each module was also modelled in a manufacturing discrete event simulator (ManPy) to evaluate serial and parallel workflows. Model accuracy was evaluated by comparing the simulator's completion times for single- and multi-patient queues to those of the RPA system. Reliable performance of the RPA was reported as number of radiation plans generated in 24 hours assuming all systems were operating. Finally, module downtime scenarios were simulated to determine their impact on baseline performance of the RPA's daily throughput. An independent t-test showed that the discrete event simulator realistically modelled mean processing times. The model estimated that the RPA could process 483 cervix plans, 255 H & N contours, or 258 H & N plans in 24 hours with all systems operating. Downtime simulations showed that cervix plan generation remained within 5% of its baseline throughput unless any given module (except for the plan/dose quality assurance module) went down for 3 hours or more. H & N contour generation remained within 5% of its baseline when downtime for either of its two contouring modules did not exceed 1 hour. H & N plan generation remained within 5% of its baseline until downtime for at least 2 of 5 available volumetric arc therapy (VMAT) optimization modules exceeded 1 hour, or any other module's downtime exceeded 3 hours. Plan calculation and report generation downtimes had <5% effect on output through 4 hours of downtime. The RPA architecture is robust to downtime of its individual modules and can provide a reliable service to clinics with limited resources.

Keywords: Automation; throughput capacity; deep learning; LMIC; treatment planning

Introduction

Half of cancer cases globally are best addressed by radiotherapy to treat, control, or palliate the disease. Eighty percent of this burden falls on low-income countries, which possess just 5% of the global radiotherapy resources for treatment [1, 2]. This shortage is most evident in staffing, as it is projected that 40,000 person-years of training are needed to address the global deficiency in medical physicists alone [1]. The Radiation Planning Assistant (RPA) is being developed to provide web-based automated contouring and treatment planning services for clinics with limited resources, including staffing shortages, to meet demand for treatment. The goal is to provide consistent, high-quality contours and plans, thereby reducing the time needed for medical physicists/treatment planners and radiation oncologists to spend preparing each patient's plan; this is a process that can take several hours or even days when performed manually [3]. For example, the RPA was designed to address staffing deficiencies with its automated workflow in countries such as Zimbabwe, which in 2015 had a total of eight radiation oncologists, five medical physicists, and 30 radiation technologists for a population of 14.5 million [1]. By contrast, the 2019 American Society for Radiation Oncology staffing model requires one therapist per 90 patients treated annually, and a minimum of one physicist and one radiation oncologist available during treatment hours [1]. Currently, the RPA is fully equipped to generate 4-field box and volumetric arc therapy (VMAT) plans for cervix, VMAT plans for head & neck, and postmastectomy chest wall plans with field-in-fields optimization. Thus, the RPA can create plans with a range of complexities. Eventually we expect it to have a portfolio of purposely developed tools such that users can choose the planning approach that is best suited to their local situation. Similarly, we will build on the current automated contour and plan verification techniques, carefully managing risk as we deploy the RPA to clinical use.

Previous studies have validated the RPA's proficiency in automated planning, time efficiency over manual planning, and the reliability of high-quality plans generated [3-8]. For example, in a study of the RPA's time efficiency, Kisling et al showed that automatically generating postmastectomy radiation therapy plans can take a mean of 38 minutes (range 28-52) per plan (for which human input takes only a few minutes), compared with manual planning, which can take upwards of 2 hours [6, 9]. That study demonstrates that leveraging automated planning through the RPA should be able to facilitate the preparation of a large volume of plans per day without the need for staff changes. In terms of quality, Olanrewaju et al found that radiation oncologists scored 88% of automatically generated VMAT plans for head and neck cancer as "use as is" [8]. In a similar study investigating the clinical acceptability of automatically generated 4-field box plans for cervical cancer, Kisling et al found that 88-100% of the field apertures were clinically acceptable [11]. Thus, these automatic tools have the potential to bring consistent, high-quality plans to the clinic with minimal user intervention. The current study now turns the focus to reliability in terms of expecting a certain output of automated plans even in unexpected situations, like system downtimes. The goal is to repeat the present study later as more disease sites are deployed to our clinical system.

The implementation of a high-throughput, high availability automated planning system for radiotherapy clinics could substantially increase efficiency [10]. Practices across all staffing and resource availability levels can greatly benefit from a system that can throughput multiple disease site plans at twice or greater the rate of human workflow, while maintaining high quality. A previous sensitivity analysis estimated that increasing efficiency with the same available radiotherapy resources, and/or reducing treatment-planning and quality-assurance time by 50%, would reduce radiotherapy facility operating costs by 21% and would yield a net monetary benefit (cost of investment – economic return) of up to $67.7 billion, thereby scaling up radiotherapy services in lower-middle-income countries [1]. Implementing automation is the prime solution to achieve higher plan throughput at lower cost. Increasing efficiency will be necessary for addressing existing shortfalls in radiotherapy access. In terms of patient numbers, we estimate, based on previous studies, that a reach rate goal of 50% of patients who currently do not have access to radiotherapy in areas served by projected RPA users in Latin America and South Africa would be about 40,000 patients. As the use of the RPA expands, a 10% global reach rate goal would be about 100,000 patients [1].

Materials and Methods

Overview of the RPA

The RPA is a fully automated treatment planning tool that designs patient-specific radiation treatments on a request basis. It can be used alongside the user's commercial treatment planning system to create 3D treatments based on planning computed tomography (CT) scans provided by the user. A summary of the user-side workflow is shown in Figure 1. The workflow was designed to maximize the number of clinics that the RPA can reach, while minimizing running cost. The RPA tools are provided through a webpage, with no local software installation. The workflow begins when the user uploads a planning CT scan set to the RPA service website and completes a service request detailing patient information. Then, the desired treatment-specific parameters and the linear accelerator are selected [3]. Once the CT set and the corresponding service request has been sent to the RPA servers, a master scheduling computer sends jobs to the respective clusters for contouring and planning. The task assignment is performed on a first-come, first-serve basis.

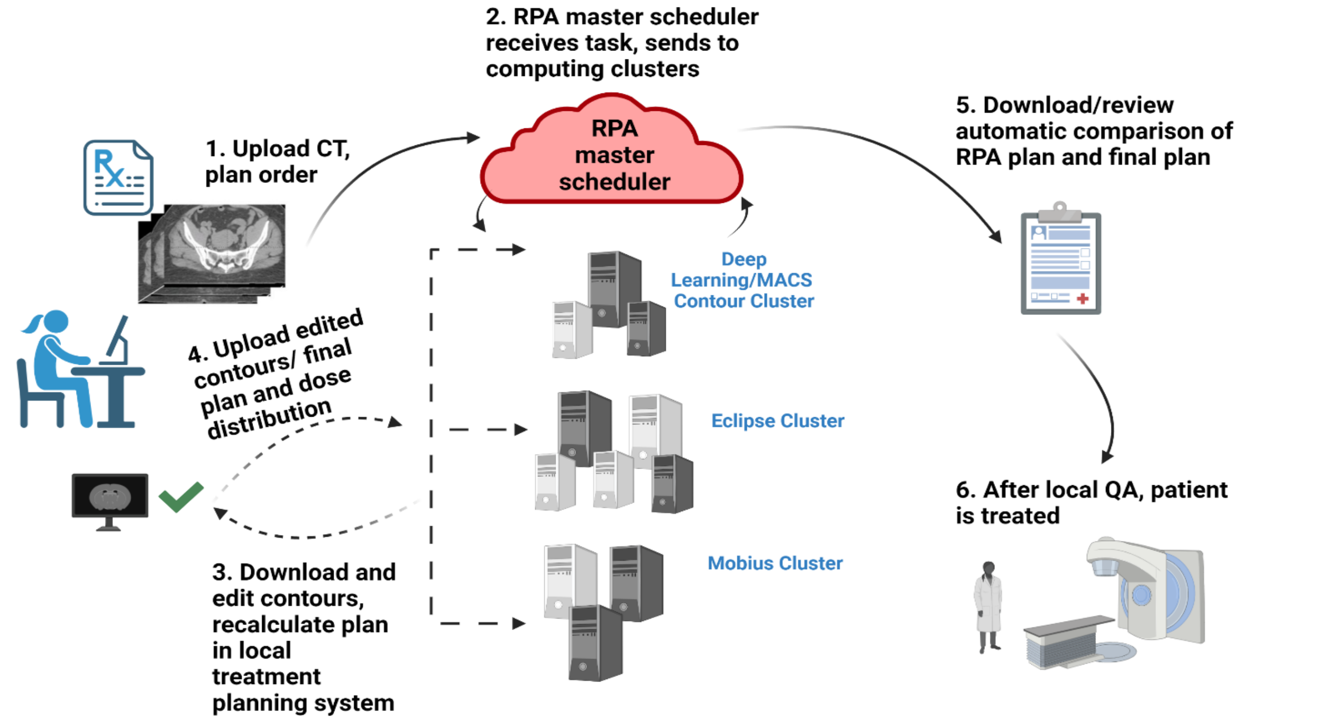

Figure 1. The user-side workflow of the RPA.

1) The user uploads a planning CT image set and completes a plan order form on the RPA website. 2) The RPA generates a set of contours and a proposed radiotherapy treatment plan. 3) Depending on the service requested, the user will receive a PDF at the time of contour generation or after plan generation to edit contours and recalculate plan in the institution's local treatment planning system. 4) The user can then re-upload the edited contours to the cloud, where the RPA will plan on the edited contours and/or run quality assurance on the RPA's final plan. 5) The user can download and review the final RPA plan, and 6) send to the local treatment planning system for patient treatment. Under the RPA master schedule, the multiple clusters involved in the contour and plan calculations are depicted. The workflow in these clusters is elaborated in Figure 2.

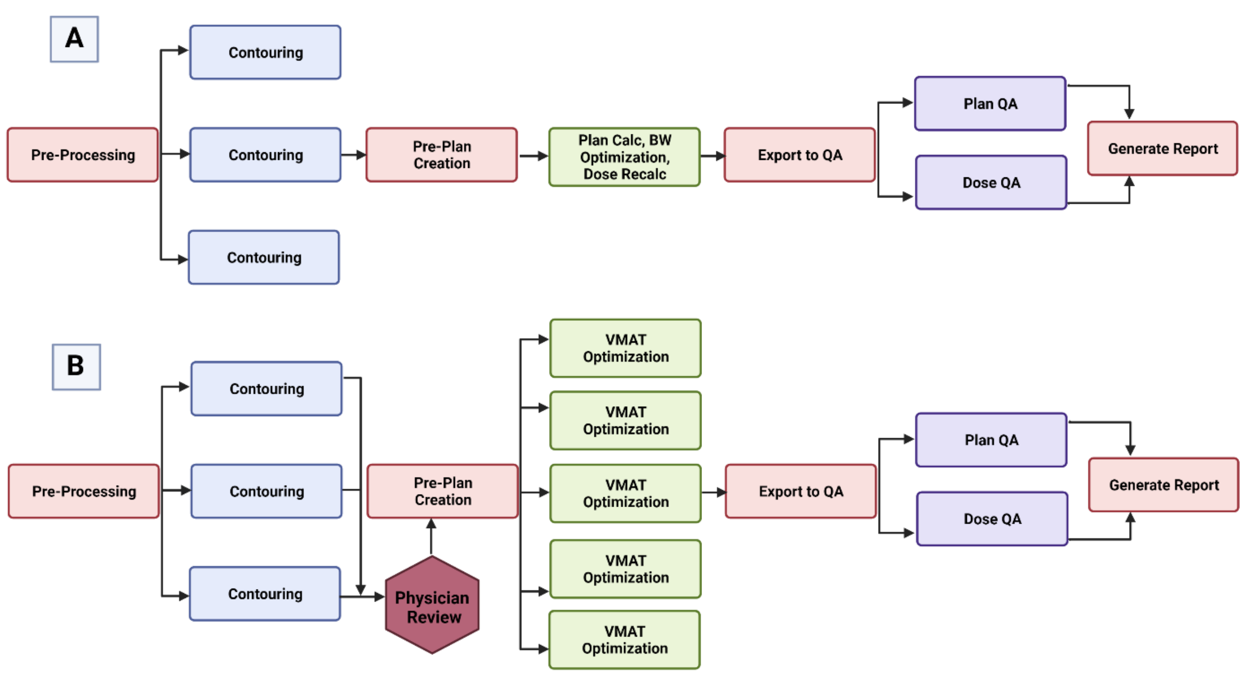

Figure 2. Workflow maps for A) cervix and B) head & neck automated planning services. Both services begin when the RPA master scheduling computer pre-processes the task and sends images to the contouring stations for parallel delineation of target volumes and normal structures.

The master computer receives contouring results and creates a pre-plan, which is sent to the dedicated treatment planning system (Eclipse) for plan calculation, beam weight (BW) optimization, dose recalculation, and/or volumetric modulated arc therapy (VMAT) optimization. The master computer once again receives these results and exports the plan to the dedicated dose verification system, Mobius, for quality assurance (QA) while Eclipse optimizes the plan in parallel. The results of plan and dose QA are re-iterated to the master scheduling computer, which generates a final report.

A summary of the preliminary automated planning workflow and its many components is detailed below.

1. The RPA master scheduling system receives the job request from the user through an internet connection and performs a series of pre-processing tasks.

2. The master scheduling computer transfers the job to the dedicated deep-learning contouring cluster (i.e. part of the central RPA system) for delineation of target and organs at risk [7, 11–13].

3. The master scheduling computer receives the generated contours and sends the job to the dedicated, in-house treatment planning system cluster, Eclipse (Varian Medical Systems, Palo Alto, CA) for dose calculation and plan optimization.

4. The final treatment plan is sent to the dedicated dose verification system, Mobius (Varian Medical Systems, Palo Alto, CA) cluster for quality assurance (QA) – specifically verification of the accuracy of the calculated dose.

5. The final Mobius plan is sent back to the master scheduling computer, and additional plan verification is performed. A final report summarizing the results generated by steps 1-3, except for Mobius results (which are used for internal monitoring of the processes), is compiled and communicated to the user as a PDF available on the website dashboard. The PDF, final contours, and plan files generated by the RPA are made available to the user to download and upload into their local treatment planning system (TPS) review. At this stage, dose recalculation and any other necessary edits can be made by the physician before delivery.

Mapping the server-side workflow

To generate RPA throughput estimates for scenarios that are not readily testable with the research version of the RPA, we built a discrete event simulator with Python software (described in the next section). We created a simplified process map of the server-side RPA workflow (Figure 2) to input into the discrete event simulator, which describes the order in which the RPA creates a radiation plan. Each task completed by the RPA was transformed into a "module" or box that represents a single process completed (contouring, plan creation, optimization, quality assurance). Therefore, the final workflow for each disease site describes all steps that different computer clusters complete serially to generate one radiation plan. Figure 2 depicts the final workflow maps composed of these modules used to build the simulation model. The "Plan QA" and "Dose QA" modules, completed by Eclipse and Mobius, occur in parallel for a single patient, and the step that finishes first varies between cases. For the sake of the model, we have simplified this process into one module with a two-patient capacity whose start and end times are determined by the first and last tasks to be received by this parallel process for the patients sampled. Notably, the head & neck planning is separated into two different workflows: "contours" and "plans." This is because the RPA offers these two as distinct services, where the client could choose to generate contours on their own and submit for plan generation, or auto-generate contours with the RPA but then must intervene to approve contours before the planning workflow begins. It is difficult to model the expected wait time between receiving RPA-generated contours and physician approval to begin the planning stage, so this was not included in the simulation.

Building a model with Manufacturing Python Simulation Software

"Manufacturing in Python," or ManPy, was used to generate simulations of the RPA workflow [14]. ManPy is a discrete event simulator that is part of a series of solutions for the defined objectives of DREAM (simulation-based application Decision support in Real-time for Efficient Agile Manufacturing). This package allows the user to create boxes or "modules" with specified metrics like the time it takes to process one part, and then connect modules like a manufacturing process. This is useful to simulate throughput systems like the RPA, whose "part" is a radiation plan. ManPy has been used as part of a free open simulation development platform in applications by European manufacturing companies, information technology consultants, and the open-source community and researchers [15]. One study demonstrated how its tools can solve a resource-constrained scheduling problem with a variable-intensity approach; another showed the effectiveness of a ManPy-based activity cycle diagram of homogenous job shops having inline cells to model an electronics fabrication line [16, 17].

We applied ManPy to simulate the RPA workflow as follows. For modeling automated planning process steps like contouring, plan creation, optimization, and quality assurance, a ManPy "machine" object, or module, was created. Each module allowed input parameters of distribution type, mean, and standard deviation of running time. Between each machine was a queue object, where 'parts' (in this case, radiation plans) wait to be processed by the next module when parallelization is not available, or when the capacity of the machine is otherwise exceeded. For example, plans that come out of the "Contouring" stage, which has 3 modules to represent the RPA's ability to contour for 3 radiation plans at a time, must wait in a queue to be processed by "Pre-Plan Creation" stage, which can only process 1 plan at a time. To model failures in the system, each machine specifies a "mean time to failure" metric; with this parameter, we modeled scheduled failures of each module. Patient plans were modeled as a "part" using the Entity class, generated either on a pre-specified timeframe or initiated instantaneously as a specified number inside a queue at the beginning of the workflow.

Sampling the time distributions

Baseline completion times for each module were measured by processing 25 4-field box cervix and 25 VMAT head and neck (H & N) patient datasets through the entire workflow. The time distribution for each task was obtained via the time logs produced by the RPA after the patient enters and leaves each module (contouring, optimization, QA, etc). To assess the ability of the model to accurately include queues in its time calculations, we simulated these tasks for single-patient scenarios and for a pre-loaded queue of each of the two sets of 25 patients.

Each module was modeled in a manufacturing discrete event simulator (ManPy) to evaluate serial and parallel workflows. Model accuracy was evaluated by comparing the simulator's completion times for single- and multi-patient queues to those of the real RPA system. Finally, many module downtime scenarios were simulated to determine their impact on baseline performance of the RPA's daily throughput.

Validating model performance with RPA

The ManPy model's output prediction was validated against the RPA by measurement of completion times in two main scenarios:

A. The distribution of times to complete one patient

B. The time to receive 25 plans back when the queue is filled instantaneously.

For scenario A, the resulting means of the two distributions verified whether the distributions were similar between the model and the research RPA via a two-tailed, heteroscedastic statistical t-test at a 95% significance level. For scenario B, we verified whether the time to get the plans back agreed with that of the RPA by direct comparison of the timing results.

Generating model estimates for throughput

After model verification was completed, we simulated several scenarios not readily testable with the research RPA:

A. Total plans completed in an 8-hour workday with no machine downtimes (assuming an infinite number of requests in the queue)

B. Plans completed in 8-hour workday when one machine fails

    a. Failure duration: 1-7 hours

    b. Failure victim: one machine in the workflow at a time

C. Plans completed in an 8-hour workday with new system resource allocations.

Results of these simulations were reported as number of radiation plans generated given the time constraint and resource constraint proposed by each scenario A-C.

Results

Completion times for single patient queues in the RPA system

The time distributions for the RPA system architecture implemented in the ManPy model for completing an individual patient plan are summarized in Table 1. The time for each module in the workflow to complete one patient plan was taken as a measured average from multiple single-queue patient runs through the RPA system. These values were used as a baseline for validating the performing of the ManPy model results. On average, cervix planning was completed in under 20 minutes; contouring and quality assurance of plans were the most time-consuming processes. H & N planning times were longer, taking about 45 minutes on average. They also shower higher variability in time distribution to complete plans, likely due to higher variability in complexity of cases. For H & N planning, VMAT optimization and quality assurance were the most time-consuming processes.

Model Validation

Results for cervix and H & N planning model validation for an individual patient are shown in Table 2. Note that the H & N workflow has been split into contours and planning for validation, as a physician intervention point after contours are generated is required to begin the planning process; this pause in the workflow is not being considered as part of the RPA-dependent time. A two-tailed t-test indicated that the means were similar between the model and the RPA. Frontloading results for 25 patients showed that both the model and the RPA converge at the same completion time for cervix planning, with an average difference of 4.98 minutes between inter-patient processing times and 2.25 minutes for H & N planning services. Similarly, minimal variation was observed between the H & N contour model and the RPA.

Table 1. A summary of the completion times for the RPA system architecture in cervix and H&N planning workflows. Shown are the mean (μ) and standard deviation (σ) in minutes for each machine's throughput of one plan

	Cervix Planning	H&N Planning
Process Name	**μ ± σ (min)**	**μ ± σ (min)**
Pre-Processing	0.19 ± 0.13	0.19 ± 0.13
Contouring	4.85 ± 0.60	9.00 ± 1.00
Pre-Plan Creation	1.83 ± 0.30	1.73 ± 0.30
Plan Calc, BW Optimization, Dose Recalc	1.60 ± 0.26	N/A
VMAT Optimization	N/A	22 ± 3.33
Export to QA	0.30 ± 0.10	0.65 ± 0.05
Plan & Dose QA	7.00 ± 0.76	10.25 ± 2.93
Generate Report	0.48 ± 0.03	0.19 ± 0.13
Total Planning Time	16.25 ± 2.18	44.01 ± 7.87

Table 2. A summary of the comparison between the ManPy model and the RPA timing. Each result is reported as the mean (μ) and standard deviation (σ) of an individual patient run through the entire workflow. A two-tailed t-test was conducted between the mean of the ManPy model and that of the RPA for each individual scenario. The mean (μ), standard deviation (σ), confidence interval (CI), and p-value of the two-tailed t-test result for each machine's throughput of 25 plans are reported.

	Cervix planning		H & N Contours		H&N Planning	
	Model	**RPA**	**Model**	**RPA**	**Model**	**RPA**
μ ± σ (min)	13.94 ± 1.35	13.75 ± 0.75	9.52 ± 0.88	9.86 ± 1.70	35.46 ± 4.62	33.93 ± 5.21
95% CI (min)	[13.38, 14,5]	[13.65, 13.85]	[9.18, 9.86]	[9.2, 10.52]	[33.65, 37.27]	[31.78, 36.08]
P, t-test	0.57		0.41		0.44	

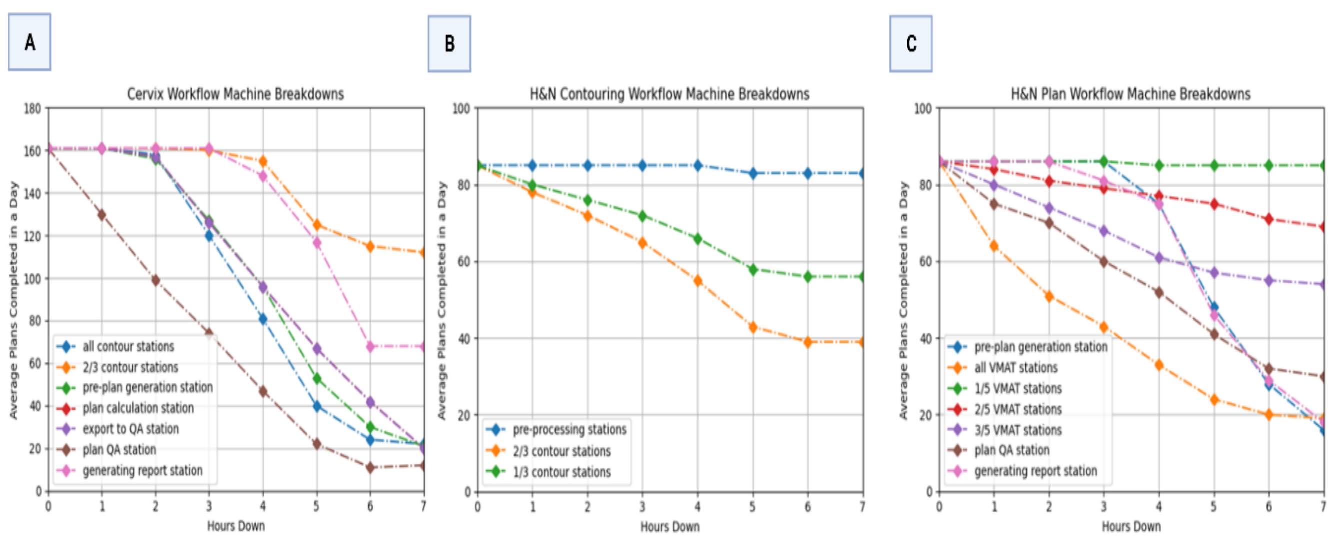

Figure 3. The average number of plans completed in a day versus the number of hours a single machine spent down.

Each machine (key) in the workflow map from Figure 3 was synthetically crashed via the ManPy software for 1-7 hours. A) Disaster behavior for the cervix automated planning service workflow. Cervix planning throughput baseline is 161 plans a day. B) Disaster behavior for the H & N automated contouring service workflow. H & N contouring throughput baseline is 85 plans a day. C) Disaster behavior for H&N planning automated service workflow. H&N planning throughput baseline is 86 plans a day. "Generating Report" and "Pre-Plan Creation" downtimes cause almost identical effects on throughput and thus are overlaid on the graph.

Throughput estimates

Eight-hour workday. The ManPy model indicated an expectation value for plans successfully processed in an 8-hour workday with no machine failures when the queue is continuously full. This excluded any plans still processing at any step at the time of simulation end. The cervix planning completed 161 plans; the H & N contours completed 85 plans; and the H & N planning completed 86 plans (assuming contours and planning requests are separate and not dependent on each other). The estimated 24-hour capacity of the RPA is thus 483 cervix plans, 255 H & N contours, or 258 H & N plans, with all resources allocated.

Disaster scenarios. We used the 8-hour workday values as a baseline performance against which to compare the throughput during a machine failure. Figure 3 summarizes the average number of plans completed in an 8-hour workday when each individual machine in the workflow was synthetically broken down for between 1 and 7 hours.

Cervix plan generation remained within 5% of its baseline throughput until any given module (except for the plan/dose quality assurance module) went down for 3+ hours. H & N contour generation remained within 5% of its baseline when downtime for either of its two contouring module's downtime did not exceed 1 hour. H & N contour generation remained within 5% of its baseline until downtime when at least 2 of 5 available VMAT optimization modules' downtime exceeded 1 hour, or when any other module's downtime exceeded 3 hours. Plan calculation and report generation downtimes had <5% effect on output through 4 hours of downtime.

Additional resource allocation. Finally, we investigated the change in 8-hour workday expectation value when additional parallelized resources were allocated to the plan workflow. For cervix planning, we found that allocating 2-3 additional contouring workstations would result in a 24-hour output of 207 plans, a 46-plan increase from the baseline with current resources. A similar increase was reached by adding 2 more "Plan & Dose QA" stations to the cervix workflow. For H & N contouring, we found a linear increase in plans completed for contouring services as contouring stations are parallelized, such that each additional contouring station would increase the throughput of contouring services. In the planning stage, a similar linear result was obtained for allocating 3-5 additional "VMAT Optimization" stations (volumetric modulated arc therapy; Eclipse).

Discussion

The purpose of this experiment was to evaluate the ability of the RPA to reliably create radiation plans, even in the presence of unexpected downtime or similar system issues. This is important as the global service demands on the RPA expand, especially if clinical users start to rely on the system to create consistent treatment plans.

In this study, the RPA was projected to be able to generate 483 cervix plans, 255 H & N contours, or 258 H & N plans in 24 hours at its current capability. This potential has profound implications for existing resource shortages. The RPA simulations showed that our automated system has the potential to quickly recover from even the improbable, long-term computing downtime scenarios posed by our experiments; we expect that these disasters would hardly inconvenience clinical workflows relying on the RPA, except for clinicians who plan on executing same-day treatments. In future studies, a tool that would allow simulation of heterogeneous case mixes (i.e., a mix of cervix H & N contour, H & N plans) being pushed to the RPA queue would prove useful for "stress testing" the performance of our auto-planning system, as well as a way to model plans being pushed to the queue from several time zones (i.e., Africa and Latin America, which contain 10 different time zones with 6–7-hour differences). Adding these complexities in a more advanced simulation tool will approximate a more clinically realistic performance metric as the RPA expands to include more disease sites for treatment planning and serves institutions around the globe.

For the RPA to successfully meet staffing shortages and improve health outcomes for its user clinics, it must create radiation plans consistently and reliably. The latter is especially important as the RPA is used to service a larger number of under-resourced clinics, increasing its throughput demand. To test its reliability, we performed a throughput capacity study of the RPA to determine the robustness of the system's response to downtimes and increasing user demand. The goal of estimating throughput under different downtime scenarios is to optimize the system's revive response; we aimed to reduce the RPA's impact on out-of-service hardware and systems on meaningful radiotherapy access in low- to middle-income countries and properly fulfill its aims [18-20]. The throughput study tested each of the multiple components of the contouring and plan generation system, which contains a variable number of parallelized computers per processes involved. We specifically tested how many plans the RPA can currently generate, the effect of different downtime scenarios on daily output, and how allocating additional resources to the RPA improves throughput capacity. The RPA's throughput response to these experimental conditions is expected to predict the ability of its current infrastructure to bring high-quality, high-throughput planning services to limited-resource clinics, minimizing their operational costs via implementing automation.

Meeting demand. The capacity of the RPA is currently greatly in excess of demand, as it can generate upwards of 100,000 plans annually and so should have the capacity needed to support radiotherapy clinics across the world. This allows the RPA to function effectively at peak demand, because that demand is not evenly spaced throughout the day. Furthermore, the RPA holds an untapped ability to generate plans to meet demands that are "pushed" from different time zones and to work the queue of plans at all hours. At the 40-hour workweek rate, we anticipate meeting demand for both our projected partner institutions in the next 3-5 years and the projected global demand in the next 10 years at a 10% reach rate for low- and middle-income countries. Additional development to add computational resources and improve the efficiency of the automated algorithms will further increase this capacity.

Response time to disaster scenarios**.** Technicians must troubleshoot machine breakdowns in a timely manner if the reliability of our system is to be sustained. To preserve cervix service throughput, we are most concerned with: upwards of 1 hour of downtime for plan & dose QA, at which point projected output declines by 30 plans; upwards of 2 hours of downtime for all contouring stations simultaneously (rare), which would cause a 30-plan decline at about 3 hours; and upwards of 2 hours of downtime of Eclipse's plan calculation service, which has a similar effect to that of having all contouring stations down. For H & N contours, the GPU responsible for contouring the tumor (labeled "fast contour", Figure 3) most affects output, decreasing it by 15 plans at 3 hours of downtime. For H & N planning, Eclipse and Mobius's parallelized plan & dose QA services or Eclipse's 5 parallel VMAT optimizers affect output the most, both individually decreasing output by about 15 plans at 2 hours of downtime.

Where automated tools provide greater radiotherapy workflow efficiency, they open opportunities for including critical steps in that workflow that are currently scarce for under-resourced clinics (such as quality assurance of contours and plans). The average correction time for automatically generated contours has been shown to be shorter than that for manually generated contours for the disease sites automated by our tool [6, 7], meaning that the speed and reliability of the RPA will enable thorough review of plans, which is vital for the safe and effective use of automation planning tools. In 2 hours, clinicians using the RPA can edit 8 automatically generated breast plans, whereas manual labor could generate only 1 breast plan in that same period. This tool is currently being implemented to temporarily mitigate the enormous staffing shortage in South Africa, one of the richest regions on the continent, which nevertheless has only 38 radiation oncologists to treat the estimated yearly incidence of 38,500 cancer cases requiring radiation therapy [1].

The goal of the RPA project is to extend its reach to offer high-quality, automatically generated contours and plans to as many clinics as possible, focusing specifically on clinics with limited resources that would find it difficult to generate these without the RPA. If successful, the RPA services must be robust and reliable. Throughput losses from complete RPA downtime over the course of 1 or more days would be virtually completely offset on the next working day of the machine, given that for current available disease sites, the RPA can generate plans at least twice as fast as plans that are manually generated [3, 5, 6, 8]. A typical center with 2 linear accelerators could reasonably onboard only 4-8 new patients a day and generate 2-5 plans per day. Hence, replacing manual plan generation with the RPA allows a more resource-efficient clinical workflow in which radiation oncology teams can plan patient schedules based on quick turnaround of RPA contours and plans.

Conclusions

In this study, the RPA was projected to be able to generate 483 cervix plans, 255 H & N contours, or 258 H & N plans in 24 hours at its current capability. Throughput losses from complete RPA downtime over the course of 1 or more days would be virtually completely offset on the next working day of the machine, given that for current available disease sites, the RPA can generate plans significantly faster than plans can be generated manually. The RPA architecture is thus robust to downtime of its individual modules and can provide a reliable service to clinics with limited resources.

Abbreviations

RPA: Radiation planning assistant; H & N: Head and neck; VMAT: Volumetric modulated arc therapy; QA: Quality assurance; TPS: Treatment planning system.

Acknowledgments

This work was supported by NIH/NCI grants UH2/UH3 CA202665 (L.E.C.) and P30CA016672 (L.E.C.). We also thank Sarah Bronson, ELS, of MD Anderson's Department of Scientific Publications, and Christine Wogan, MS, ELS, of MD Anderson's Division of Radiation Oncology, for editing this article.

Author Contributions

M. B., G. S., D. R., Z. L., E. B. D., H. K., C. E. C. and C. E. L. contributed equally to this study. All authors gave their final approval.

Competing Interests

The authors have declared that no competing interest exists.

References

[1] Atun, R., Jaffray, D. A., Barton, M. B., Bray, F., Baumann, M., Vikram, B., Hanna, T. P., Knaul, F. M., Lievens, Y., Lui, T. Y. & Gospodarowicz, M. (2015). Expanding global access to radiotherapy. *The lancet oncology*, *16*(10), 1153-1186.

[2] Datta, N. R., Rogers, S., & Bodis, S. (2019). Challenges and opportunities to realize "The 2030 Agenda for Sustainable Development" by the United Nations: Implications for radiation therapy infrastructure in low-and middle-income countries. *International Journal of Radiation Oncology* Biology* Physics*, *105*(5), 918-933.

[3] Kisling, K., McCarroll, R., Zhang, L., Yang, J., Simonds, H., Du Toit, M., Trauernicht, C., Burger, H., Parkes, J., Mejia, M. & Beadle, B. (2018). Radiation planning assistant-a streamlined, fully automated radiotherapy treatment planning system. *JoVE (Journal of Visualized Experiments)*, (134), e57411.

[4] Kisling, K., Johnson, J. L., Simonds, H., Zhang, L., Jhingran, A., Beadle, B. M., Burger, H., du Toit, M., Joubert, N., Makufa, R. & Court, L. (2019). A risk assessment of automated treatment planning and recommendations for clinical deployment. *Medical physics*, *46*(6), 2567-2574.

[5] Kisling, K., Zhang, L., Simonds, H., Fakie, N., Yang, J., McCarroll, R., Balter, P., Burger, H., Bogler, O., Howell, R. & Court, L. (2019). Fully automatic treatment planning for external-beam radiation therapy of locally advanced cervical cancer: a tool for low-resource clinics. *Journal of Global Oncology*, *5*, 1-9.

[6] Kisling, K., Zhang, L., Shaitelman, S. F., Anderson, D., Thebe, T., Yang, J., Balter, P. A., Howell, R. M., Jhingran, A., Schmeler, K. & Court, L. (2019). Automated treatment planning of postmastectomy radiotherapy. *Medical physics*, *46*(9), 3767-3775.

[7] Rhee, D. J., Jhingran, A., Rigaud, B., Netherton, T., Cardenas, C. E., Zhang, L., Vedam, S., Kry, S., Brock, K. K., Shaw, W. & Court, L. E. (2020). Automatic contouring system for cervical cancer using convolutional neural networks. *Medical physics*, *47*(11), 5648-5658.

[8] Olanrewaju, A., Court, L. E., Zhang, L., Naidoo, K., Burger, H., Dalvie, S., Wetter, J., Parkes, J., Trauernicht, C. J., McCarroll, R. E. & Beadle, B. M. (2021). Clinical Acceptability of Automated Radiation Treatment Planning for Head and Neck Cancer Using the Radiation Planning Assistant. *Practical radiation oncology*, *11*(3), 177-184.

[9] Yang, J., Amini, A., Williamson, R., Zhang, L., Zhang, Y., Komaki, R., Liao, Z., Cox, J., Welsh, J., Court, L. & Dong, L. (2013). Automatic contouring of brachial plexus using a multi-atlas approach for lung cancer radiation therapy. *Practical radiation oncology*, *3*(4), e139-e147.

[10] Rhee, D. J., Jhingran, A., Kisling, K., Cardenas, C., Simonds, H., & Court, L. (2020, October). Automated radiation treatment planning for cervical cancer. In *Seminars in radiation oncology* (Vol. 30, No. 4, pp. 340-347). WB Saunders.

[11] Kisling, K., Cardenas, C., Anderson, B. M., Zhang, L., Jhingran, A., Simonds, H., Balter, P., Howell, R. M., Schmeler, K., Beadle, B. M. & Court, L. (2020). Automatic verification of beam apertures for cervical cancer radiation therapy. *Practical radiation oncology*, *10*(5), e415-e424.

[12] Cardenas, C. E., Beadle, B. M., Garden, A. S., Skinner, H. D., Yang, J., Rhee, D. J., McCarroll, R. E., Netherton, T. J., Gay, S. S., Zhang, L. & Court, L. E. (2021). Generating high-quality lymph node clinical target volumes for head and neck cancer radiation therapy using a fully automated deep learning-based approach. *International Journal of Radiation Oncology* Biology* Physics*, *109*(3), 801-812.

[13] Rhee, D. J., Cardenas, C. E., Elhalawani, H., McCarroll, R., Zhang, L., Yang, J., Garden, A. S., Peterson, C. B., Beadle, B. M. & Court, L. E. (2019). Automatic detection of contouring errors using convolutional neural networks. *Medical physics*, *46*(11), 5086-5097.

[14] Dagkakis, G., Heavey, C., Robin, S., & Perrin, J. (2013, September). ManPy: An open-source layer of DES manufacturing objects implemented in SimPy. In *2013 8th EUROSIM Congress on Modelling and Simulation* (pp. 357-363). IEEE.

[15] Lentes, J., Eckstein, H., Constantinescu, C., Lanning, I., Rotondo, A., Holland, R., Robin, S., Bayer, M., Eichert, J., Balkan, O., Lange, O., Heilmayer, S. and Sun, C. (2015) DREAM simulation-based application Decision support in Real-time for Efficient Agile Manufacturing, DREAM Consortium Report.

[16] Olaitan, O., Young, P., & Geraghty, J. (2015, July). Variable intensity RCPSP approach to a case study flow shop. In *Proceedings of the Conference on Summer Computer Simulation* (pp. 1-9).

[17] Kim, H., Choi, B. K., & Shin, H. (2015, September). ACD Modeling of Homogeneous Job Shops Having Inline Cells. In *IFIP International Conference on Advances in Production Management Systems* (pp. 340-347). Springer, Cham.

[18] McCarroll, R., Youssef, B., Beadle, B., Bojador, M., Cardan, R., Famiglietti, R., Followill, D., Ibbott, G., Jhingran, A., Trauernicht, C. & Court, L. (2017). Model for estimating power and downtime effects on teletherapy units in low-resource settings. *Journal of global oncology*, *3*(5), 563-571.

[19] Elmore, S. N., Sethi, R. V., Kavuma, A., & Kanyike, D. M. (2017). Broken machines or broken systems: The road to meaningful global radiotherapy access. *Journal of Global Oncology*, *3*(5), 438-440.

[20] Aruah, S. C., Asogwa, O. C., Ubah, F. I., Maurice, N. N., Oyesegun, R., Ige, T. A., Coleman, C. N., Dosanjh, M. & Pistenmaa, D. (2020). Overcoming challenges in providing radiation therapy to patients with cancer in Nigeria and experience in the National Hospital Abuja, Nigeria. *JCO Global Oncology*, *6, 1232–1236*.