Medical Physics Calculations in the Cloud: A New Paradigm for Clinical Computing

Roy W. Keyes[1], Christian Romano[2], Dorian Arnold[2], and Shuang (Sean) Luan[2]

University of New Mexico, Albuquerque, New Mexico, USA

1: Department of Physics and Astronomy; 2: Department of Computer Science

UNM Particle Therapy Group

Web: cs.unm.edu/~compmed/PTG

THE UNIVERSITY of NEW MEXICO

Download this poster with your phone!

http://cs.unm.edu/~compmed/p/1/

The big picture: Cloud computing can enable routine clinical use of calculations currently deemed too resource intensive, in particular Monte Carlo. This is possible because cloud services offer on-demand, scalable, virtual clusters with pay-as-you-go pricing as cheap as ≈$0.029/(GHz◦hr)[1].

What is cloud computing?: Cloud computing is a term loosely used to describe services offering CPU time, storage, and software over the Internet. The most common features of cloud computing services are transparent, on-demand scaling of resources and pay-as-you-go pricing. Companies offering could services include Google, Amazon, and Microsoft.

The Cloud Calculation Process

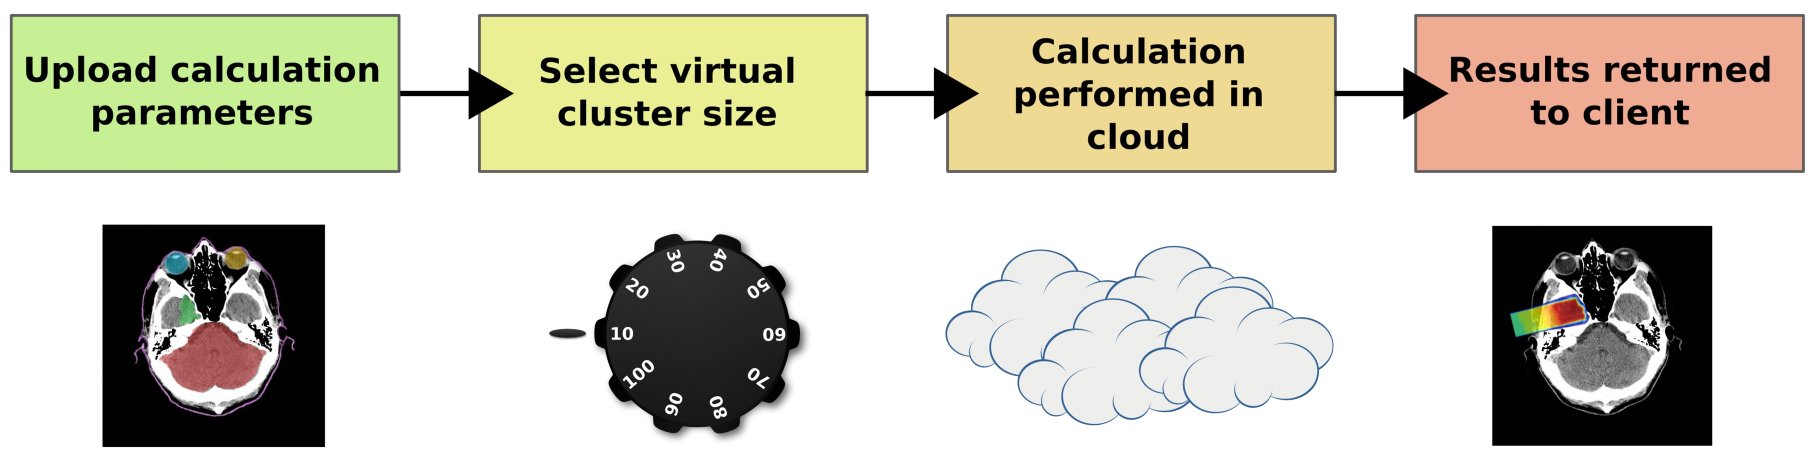

Figure 1: The cloud calculation process from the end user perspective.

Demonstrating cloud computing for medical physics: We designed and implemented a custom distributed processing framework for Monte Carlo, called *flsshd*, using Amazon's EC2 service as a virtual cluster. Figures 1 & 2 illustrate the steps in the calculation process.

flsshd = Python + bash + boto + Fluka + SSH

Coding — Amazon cloud access via Python — Monte Carlo — Networking

Proof-of-concept calculations included depth-dose curves of ^{60}Co, 10 MeV electrons, and proton beams of various energies, as well as a simple broad-beam proton plan with a CT-based voxel phantom[3].

Performance tests were made on virtual clusters of 1-200 nodes by calculating 75 and 200 MeV proton depth-dose profiles. The calculation times were modeled with Equation 1. The fit was excellent ($R^2 > 0.99999$). The linear overhead term, βn, is thought to be due to the serial nature of how flsshd initializes the calculations on the nodes. A parallelized version is being implemented and a subsequent performance increase is expected.

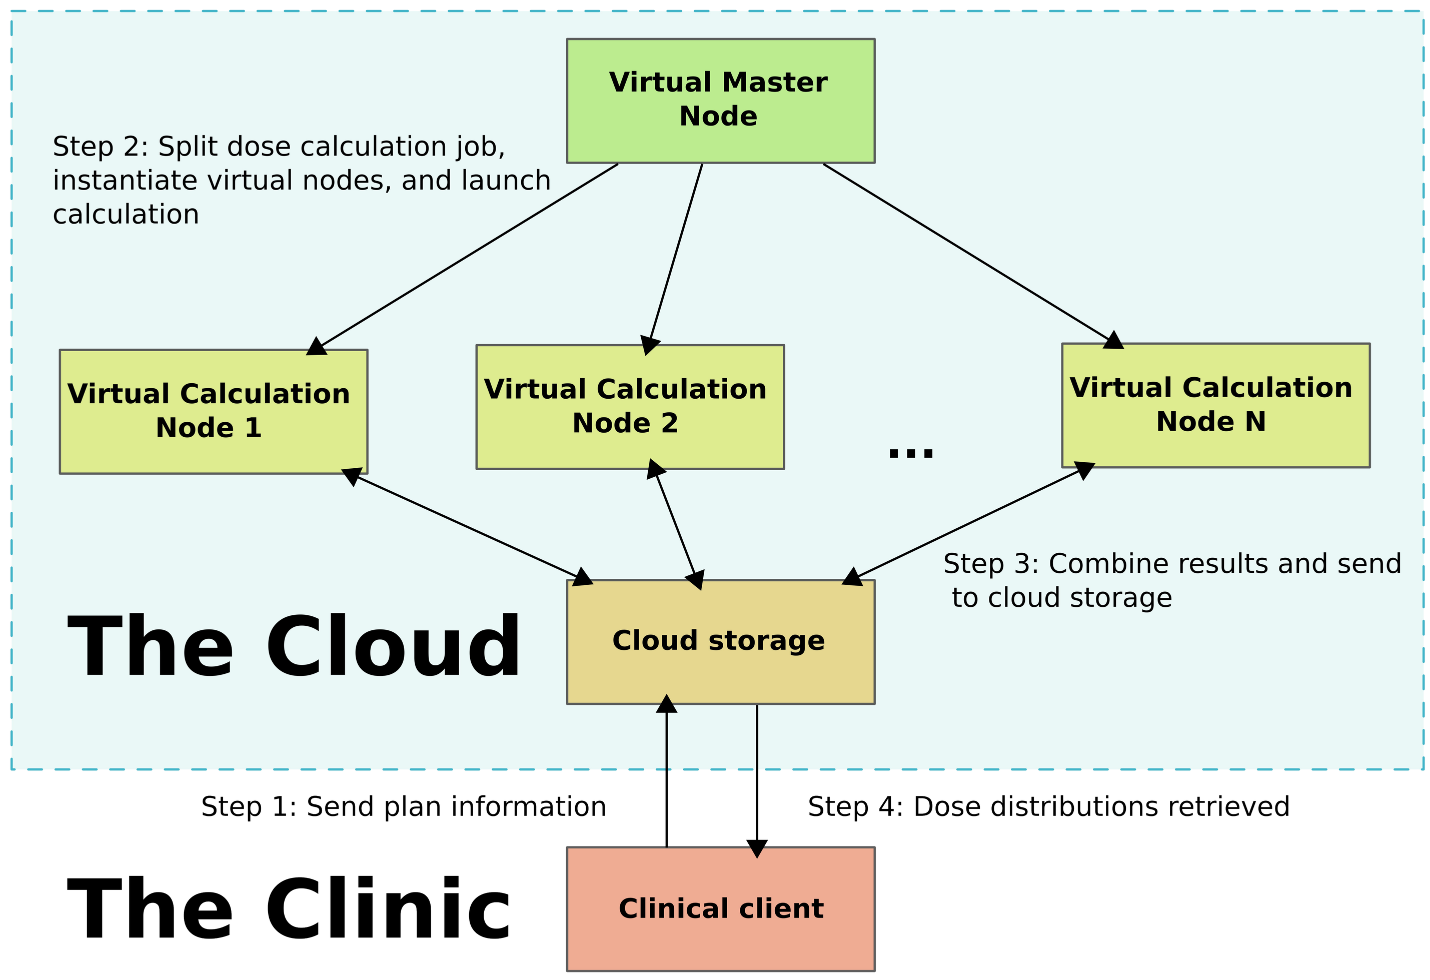

Figure 2: The cloud computing model for clinical calculations.

$$T(n) = \frac{\alpha p}{n} + \beta n + \gamma$$

Equation 1

T	Total run time
n	Number of nodes
p	Number of primary particles
α	Calculation time per primary
β	Linear overhead per node
γ	Constant overhead

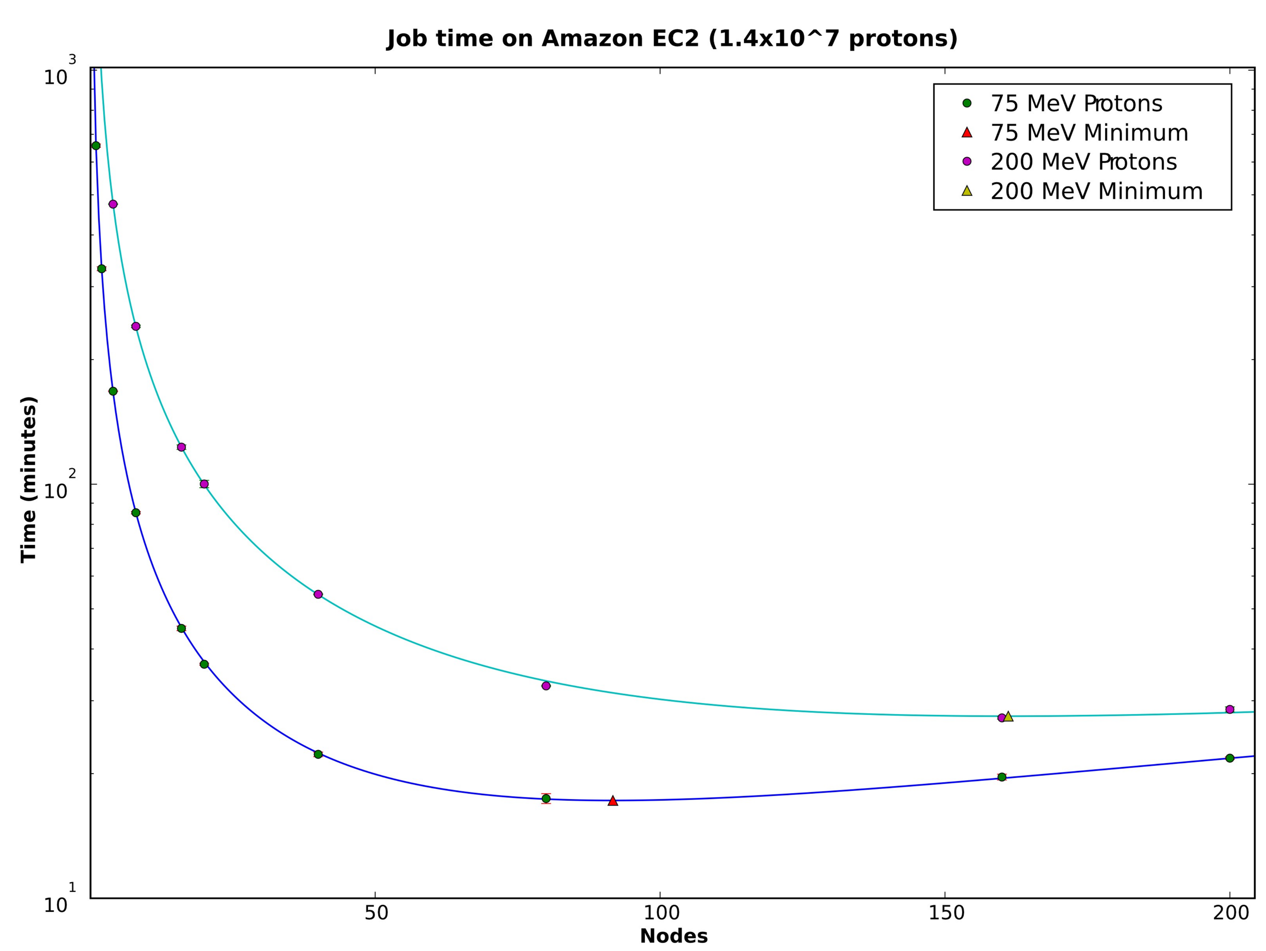

Figure 3: Calculation times for proton beam depth-dose profiles in water. The fit curves are from Equation 1.

[1] Amazon Elastic Compute Cloud 2010 http://aws.amazon.com/ec2/
[2] Fasso A, Ferrari A, Ranft J & Sala P 2005 Fluka: a multi-particle transport code. Technical report, CERN-2005-10, INFN/TC 05/11, SLAC-R-773.
[3] I. G. Zubal, C. R. Harrell, E. O. Smith, A. L. Smith, Two dedicated software,voxel-based, anthropomorphic (torso and head) phantoms, http://noodle.med.yale.edu/zubal/

Our proof-of-concept tests have been successful: We were able to remotely perform several clinically relevant calculations with an on-demand, virtual Monte Carlo cluster. Our current implementation of flsshd is limited by the linear time it takes to initialize calculations on each node. This leads to greater minimum possible calculations times at present than desired.

Cost estimates for in-house clusters are usually set around $1000/node plus $200/node per year in maintenance (a minimum cost of $53k/year over a 3 year cluster lifespan for a 100 node cluster). For a clinic serving 1000 patients per year at 100 CPU hours per patient, the price of running on a cloud service could be as low as $8,700/year in CPU time (using 3GHz CPU hours at ≈$0.029/(GHz◦hr)).

Patient data privacy can be assured by anonymizing plan data before it is encrypted and sent to the cloud for calculations.

Cloud vs. Grid: Cloud services have the advantages of ubiquitous access, lower entry barriers, orders of magnitude greater economies of scale, and backing by the largest Internet companies.

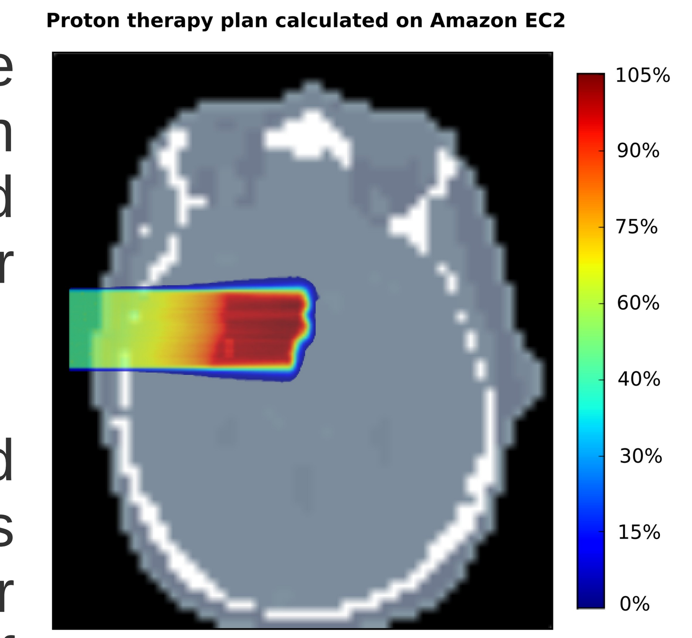

Figure 4: A simple proton therapy plan calculated on 130 EC2 nodes using a voxel phantom.

A promising outlook for clinical Cloud Computing: Cloud computing promises ubiquitous access to super computing level resources for medical physics. This has many implications for the future of radiation therapy planning and delivery. The pay-as-you-go model will allow users to replace or supplement existing infrastructure without incurring up-front costs.

Our future research plans include improving the performance of our current distributed cloud calculation framework, testing several other established distributed frameworks, optimizing costs, and developing a web portal for job submission by outside researchers.

Acknowledgments: This work was supported in part by the National Science Foundation under grant CBET-0853157.