


Imaging particle beams for cancer treatment

Jerimy C. Polf and Katia Parodi



Proton and carbon-ion radiotherapy are powerful tools for killing tumor cells, but only if the particles deposit their energy where they're supposed to.

In 2014 approximately 1 in 7 deaths worldwide were due to cancer and an estimated 14 million new cases of cancer were diagnosed. Many cancer patients receive radiotherapy either on its own or in conjunction with chemotherapy or surgery. Radiotherapy works as a cancer treatment by depositing energy through atomic and nuclear interactions in patient tissues and thereby damaging tumor cells. That energy deposition, known as the treatment dose, is measured in units of joules per kilogram of tissue, or grays. The goal is to deliver the prescribed radiation treatment dose to the entire tumor volume while minimizing or eliminating the dose received by healthy tissues and organs. Toward that end, the past 20 years have seen the development and deployment of sophisticated new treatment techniques designed to precisely target and deliver radiation to the tumor volume. (See the article by Arthur Boyer, Michael Goitein, Antony Lomax, and Eros Pedroni, *PHYSICS TODAY*, September 2002, page 34.)

In particular, the prevalence of radiotherapy

based on proton and carbon-ion beams has rapidly increased over the past 10–15 years. The distinct clinical advantage that ion beams provide over x rays was first pointed out in 1946 by Robert Wilson.¹ To first order, the rate at which proton and carbon-ion beams deposit dose in a medium is inversely proportional to the particles' kinetic energy. As a result, the dose delivery rate is lowest when the beam first enters the patient, gradually increases with depth as the particles lose energy, and culminates in a localized sharp increase, known as the Bragg peak, just before the beam stops. The depth of the sharp dose falloff just beyond the Bragg peak, called the beam range, is a function of the proton or ion energy used for treatment. By carefully selecting and modulating the beam energy, radiation oncologists can choose the beam range so that the high-dose Bragg peak is precisely delivered to the tumor while critical organs beyond the tumor are almost entirely spared. The ability to deliver more dose to the tumor and less to the surrounding healthy tissue means, in principle, that patients are less likely to experience posttreatment complications and side effects and are more likely to be cured of their cancer.

Despite the promise and potential of the Bragg



Jerimy Polf is an assistant professor of radiation oncology at the University of Maryland School of Medicine in Baltimore. **Katia Parodi** is a professor and chair of medical physics at Ludwig-Maximilians University in Munich.

peak, the ability to fully exploit the advantages of proton and carbon-ion beam therapy to treat patients is still limited. There is a degree of error in any prediction of location of the Bragg peak and the beam range. That beam-range uncertainty may cause the Bragg peak to overshoot or undershoot the tumor and damage adjacent healthy tissues.

One promising strategy to limit beam-range uncertainty has been the development of techniques to image the beam as it passes through the patient during treatment. Methods to image the beam *in vivo* would allow clinicians to verify that the treatment is being delivered to the tumor as intended and that the tumor is receiving the full prescribed radiation dose. Recent *in vivo* imaging research has exploited the basic physical mechanisms through which particle radiation interacts with matter to produce low-pressure acoustic signals and high-energy gamma rays. Detecting and imaging those secondary emissions outside the patient can provide a real-time visualization of the treatment beams inside the patient.

The perils of uncertainty

The causes of the beam-range uncertainty can be divided into two main categories: treatment planning and treatment delivery. Before a patient begins a course of radiotherapy, a treatment plan must be developed based on a specialized computerized tomography (CT) scan known as a simulation. From the CT simulation (possibly coupled with other diagnostic images), clinicians determine the tumor location and plan the parameters of daily treatment. The plan includes the number of beams (typically 2–4), the directions at which they're oriented, and the dose they're intended to deliver. Part of the planning process includes calculating both the exact range of each treatment beam needed to cover the entire tumor and the beam energy needed to produce that range. However, the accuracy of those calculations is compromised by noise and distortions in the CT image, uncertainty in the composition and density of patient tissues, and limitations in the algorithms used to determine the beams' dose delivery rates in different tissues.

The second category of range uncertainty arises during delivery of the treatment. For a treatment plan to be effective, the patient must be immobilized and positioned on the treatment couch in a way that matches the CT simulation. The positioning and alignment for each treatment session is done using a robotic couch and onboard x-ray or CT imaging systems to ensure that the patient's bone and soft-tissue landmarks are aligned as they were at the time of treatment planning. However, exact reproduction of the planned patient alignment is not possible, for several reasons. First, the onboard imaging systems are limited in contrast and resolution. Second, over the course of radiotherapy, which can last 30 days or more, patient anatomy can change due to weight loss, tumor shrinkage, or normal tissue swelling.




Figure 1. For a lung tumor (outlined in green) abutting the heart (pink), the ideal treatment plan (a) would use a single proton beam (outlined by dashed white lines) that stops at the deepest edge of the tumor. Due to the beam-range uncertainty, however, a margin (orange shaded area) must be added to the treatment area targeted for the full prescribed radiation dose. The end of the beam range is thus inside the heart, which may suffer severe damage or functional complications. To avoid such risk to critical organs, a suboptimal plan (b) with two beams might be used instead, even though it now delivers a low to intermediate radiation dose to the healthy lung.

Finally, patients may wiggle, cough, scratch an itch, or otherwise shift position during the several minutes between the end of the alignment process and the conclusion of the daily treatment.

To mitigate the effects of beam-range uncertainty and ensure that the entire tumor receives the prescribed radiation dose, an additional thickness of tissue around the tumor, the range-uncertainty margin, is included in the target volume that receives the full treatment dose. The range-uncertainty margin is typically chosen to be 2 mm plus 3.5% of the beam range, so for tumors deep in the body, 1 cm or more of tissue might be added to the treated target volume.²

As an example, consider a patient being treated for a lung tumor adjacent to the heart, as shown in figure 1. If there were no uncertainty in the beam range, the ideal treatment plan would be to use a single beam incident from the side and stopping at the deepest edge of the tumor adjacent to the heart, as shown in figure 1a. However, that arrangement would never be used, because the beam-range-uncertainty margin would include part of the heart, and the full radiation dose may do severe damage to that critical organ. Instead, nonideal beam arrangements such as the one in figure 1b are routinely used to avoid shooting the beam directly at the heart. As a result, the tumor can be fully treated and the heart is spared from high doses, although one of the treatment beams now passes through the lung. The low to intermediate doses thereby delivered to the healthy lung tissues can also have long-term implications for the patient's health. But the patient has two lungs and only one heart, so the risk of lung damage is regarded as more acceptable.

Indeed, adding the range-uncertainty margin and using suboptimal beam arrangements often involves giving away the advantage of the Bragg peak—the sparing of healthy tissue beyond the beam range. When healthy tissues and organs in the range-uncertainty margin are subjected to high doses of radiation on purpose, the patient can experience harmful and potentially life-altering complications and side effects.




Figure 2. Three types of secondary emission occur when a proton or ion treatment beam interacts with tissue. In thermoacoustic emission, the particle beam interacts with atomic electrons, as shown in the top panel, and locally heats the tissue to create a pressure wave, as shown in the bottom panel. Positron-annihilation gammas are created through inelastic nuclear interactions that leave behind positron-emitting isotopes; each emitted positron annihilates with an electron to produce two 511-keV gamma rays. Prompt gammas are produced when nuclear scatter events promote tissue nuclei to excited states that decay through the emission of single gammas, whose energies depend on the element involved.

Secondary emissions

To address the critical need to reduce or eliminate the uncertainty in the delivered beam range, a broad community of physicists and engineers has worked over the past 15 years to develop ways to measure and assess proton beams during treatment. This article focuses on one aspect of that effort, the imaging of secondary emissions created as the beams pass through the patient. As shown in figure 2, secondary emission can take the form of either thermoacoustic waves that arise from electromagnetic interactions between the beam and the tissue or gamma rays created through nuclear interactions. Both thermoacoustic waves and secondary gammas are directly correlated to the delivery of dose by the treatment beam, so imaging when and where they're emitted provides a path to verifying the beam range.

The most prevalent mechanism of dose deposition by proton and ion beams is through the electromagnetic transfer of energy to atomic electrons. In what has recently been termed the ionoacoustic process, the electromagnetic interactions produce a local increase in temperature, and the resulting thermal expansion launches a thermoacoustic pressure wave.³ For a given amount of deposited dose, the amplitude and frequency of the ionoacoustic signal strongly depend on the sharpness of the energy deposition in space and time; both quantities are maximized by an intense and highly localized energy deposition—such as the Bragg peak—delivered over a short time.⁴ From the speed of sound in tissue and the time of flight of the thermoacoustic wave to one or more transducers, it's possible to calculate the Bragg peak position, as shown in figure 3a.

Nuclear interactions between charged-particle beams and tissue produce secondary gamma rays through two distinct processes. First, inelastic interactions between protons or ions and tissue nuclei

can produce short-lived radioisotopes, such as carbon-11 and oxygen-15, that decay via positron emission. The emitted positrons annihilate with electrons to produce pairs of simultaneous 511-keV gamma rays emitted in approximately opposite directions. Those positron-annihilation gamma rays can be imaged using positron emission tomography (PET) to monitor beam delivery *in vivo*.⁵ Second, nuclear scattering processes can leave tissue nuclei in excited states. When the excited nuclei then decay to their ground states within nanoseconds, they emit single photons known as prompt gamma rays.⁶ Because each element emits prompt gamma rays with a unique spectrum of energies governed by its nuclear energy levels,

prompt-gamma imaging also enables spectroscopic analysis of irradiated tissue. The nuclear processes that produce both positron-annihilation gammas and prompt gammas occur only where the beam is interacting with patient tissue; the resulting distributions of gamma emission are well correlated with the dose deposited by the treatment beam, as shown in figures 3b–d; figure 4 shows images of prompt-gamma and positron-annihilation emission reconstructed from measurements made during and after delivery of a proton treatment beam to simulated patients consisting of tanks of water and gelatin.

In vivo imaging

Of the three forms of secondary emission, positron-annihilation gamma rays have been the most investigated in clinical settings so far,⁵ because they can be imaged using existing PET scanners. Despite a few pioneering developments of dedicated instrumentation for direct installation onto the beam-delivery system (so-called in-beam PET), most clinical trials have relied on conventional PET instrumentation originally developed for the purpose of diagnostic nuclear medicine imaging, with the imagers installed either inside the treatment room (in-room) or outside the room (off-line). Clinical tests have shown that PET imaging can identify inaccuracies in the beam range for tumors that are at high risk of treatment delivery error, including deep-seated tumors requiring a long beam range, such as those in the abdomen or pelvis, and tumors located among many interfaces between bone, soft tissue, and air cavities that complicate beam propagation, such as those in the head or neck. Studies conducted on tumors in those regions all suggest that in-beam PET could prove highly beneficial for determining the delivered beam range, guiding refinements of




Figure 3. Secondary emission from cancer patients illustrates the promise of *in vivo* imaging. (a) A liver cancer patient was treated with a single proton beam incident in the direction of the thin red arrow. The resulting complex acoustic pulse, shown by the green profile, was detected by a hydrophone in the position of the white arrow. With its time scale converted to distance, the acoustic pulse conveys information about the beam penetration depth in the patient. (Adapted from ref. 14.) (b) A prostate cancer patient, also treated with a single proton beam, received the estimated radiation dose shown by the color scale. Simulations of the corresponding distributions of positron-annihilation gammas (c) and prompt gammas (d) both correlate well with the distribution of dose. The lower energy threshold for prompt-gamma production results in a signal that extends much closer to the end of the beam range, indicated by the yellow arrows. (Panels b–d adapted from M. Moteabbed, S España, H. Paganetti, *Phys. Med. Biol.* **56**, 1063, 2011.)

beam-range calculations, and adapting the treatment plan over the course of radiotherapy.

Conventional PET scanners, on the other hand, are tailored for diagnostic imaging, so their limited sensitivity and performance present major challenges for *in vivo* range monitoring. The concentrations of radioactive nuclei are typically several orders of magnitude lower in beam therapy than in diagnostic PET imaging, so low counting statistics become a problem. In-beam imaging systems can gather data during or shortly after irradiation and thus capture a large portion of the emission from the most important positron-emitting radioisotope, ^{15}O , whose half-life is two minutes. When the patient must be moved between treatment delivery and PET scanning—especially to an off-line PET system in a different room—the time delay means that much of the signal is lost. Furthermore, physiological processes such as blood flow cause some of the positron emitters to diffuse away from the treatment area over time and thus spoil the physical correlation between treatment delivery and measured activity. Researchers are working to develop both modern in-beam PET instrumentation tailored to

in vivo beam-range verification^{7,8} and computer-assisted tools for adjusting the treatment plan when PET detects an offset between the delivered and intended beam ranges.

Imaging of prompt gammas could overcome the major shortcomings of PET monitoring because prompt gammas are not affected by physiological processes and their production cross sections are much more favorable. Initial studies have found a good correlation between dose delivery and prompt-gamma emission,^{9,10} and small shifts in the Bragg peak position can be detected by measuring prompt-gamma emission during treatment.^{11,12} Recent studies have also shown that the intensities of the characteristic prompt-gamma spectral lines from individual elements are directly proportional to the concentration of the elements in tissue,¹³ so prompt-gamma imaging could provide a path to spectroscopic analysis of irradiated tumors and healthy tissues.

Because prompt gammas are relatively high in energy (2–15 MeV), *in vivo* imaging has thus far been a challenge. Existing imaging systems, designed for gamma energies of a few hundred keV, have poor




Figure 4. Simulated patients—tanks of water or gelatin—are irradiated with proton beams for proof-of-principle gamma imaging. In each case, the dose distribution in the upper panel is calculated, and the gamma distribution in the lower panel is measured. (a) A 150-MeV clinical proton pencil beam impinges on a water tank, and the prompt-gamma (PG) emission is imaged with an experimental Compton camera specially designed for the purpose. (b) Positron-annihilation gamma emission from a tank of tissue-like gelatin is imaged with a commercial diagnostic positron emission tomography instrument shortly after irradiation with a 177-MeV proton beam. The strong signal near the entrance region in the PET image is a result of carbon-11 activation in the carbon-rich walls of the tank. (Panel b courtesy of Julia Bauer, Heidelberg Ion-Beam Therapy Center.)




Figure 5. Patient outcomes can be greatly improved by *in vivo* imaging that reduces beam-range uncertainty by just a few millimeters. In this breast cancer patient, the intended treatment volume is indicated by the purple line, and the direction of the treatment beam is shown by the blue arrow. **(a)** To account for the beam-range uncertainty, the treatment volume must be expanded to the region enclosed in yellow. As a result, significant dose (shown here as a percentage of the prescribed dose) is delivered to the heart (red line) and the left anterior descending coronary artery (LAD; white line). **(b)** With range-verified proton radiotherapy, no expansion to the treatment volume would be needed. Dose delivered to the heart and the LAD is greatly reduced, as is the patient's risk of radiation-induced heart disease.

detection efficiency in the 2–15 MeV range, and their mechanical collimators are ineffective at those energies. Many scientists and engineers are working to design and build detectors optimized for prompt-gamma imaging. Systems under development include so-called knife-edge and multislit collimators, integrating scintillation detectors, and multistage Compton cameras. The goal is to measure an adequate number of collimated prompt gammas during a single treatment session to produce one-dimensional profiles and 2D or 3D images of the prompt-gamma emission from the patient. Work is also under way to develop software to overlay prompt-gamma and CT images for visual inspection to confirm that the treatment is being delivered as intended.

Thermoacoustic emissions were first investigated clinically in Japan in the 1990s, when a liver cancer patient was being treated using a specially designed pulsed proton accelerator.¹⁴ However, that proof-of-principle measurement was limited by acoustic instrumentation that was not yet optimized for the amplitudes and frequencies required by the new application. Furthermore, first-generation proton-beam treatment systems produced broad beams with complicated timing structures that resulted in highly complex ionoacoustic signals.

Recently, though, ionoacoustics has seen a renewed interest thanks to new-generation beam-therapy systems that use compact pulsed accelerators and monoenergetic narrow pencil beams, which produce ionoacoustic signals with much more favorable signal strengths and time structures.⁴ Still, the induced thermoacoustic waves are greatly attenuated in tissue, and it's not always possible to find accessible spots on the patient's skin to place the acoustic transducers. Those constraints probably preclude the use of ionoacoustic imaging for cer-

tain treatment sites such as the upper head, due to limited transmission through the skull. But the technique is applicable for many cancers, such as those in the prostate, liver, and breast, that are commonly treated with external-beam radiotherapy and challenged by the problem of range uncertainty. Ionoacoustic images of the Bragg peak could be combined with conventional ultrasound images of internal anatomy to confirm the beam range in the patient.

Importantly, ionoacoustic imaging, unlike gamma-ray imaging, derives its signal from electromagnetic interactions, which are also the dominant mechanism by which the beam transfers energy to tissue. The ionoacoustic signal may thus be more closely linked to the actual dose deposition in the patient than the secondary-gamma signal, which depends on nuclear interactions that are responsible for only a small fraction of the dose delivered.

Improving treatment outcomes

Emerging methods for imaging the path of a proton or ion beam through the patient could greatly reduce the uncertainty in locating the Bragg peak and thereby reduce the need to add uncertainty margins around the tumor to ensure that it receives the full prescribed radiation dose. The question then becomes, how would that development improve the final outcome of radiotherapy for the patient?

As an example, figure 5 illustrates the distribution of dose delivered to a patient receiving proton radiotherapy for breast cancer. For a standard treatment (figure 5a), the volume that receives the full treatment dose, enclosed by the yellow line, includes the target volume to be treated for cancer (enclosed by the purple line) plus the extra margin to account for beam-range uncertainty. As figure 5b shows, the dose delivered to the heart and lung can be significantly decreased if the range uncertainty margin

can be reduced or even eliminated through daily *in vivo* range imaging. In that case, only the desired treatment volume would receive the full treatment dose. Of particular importance is a large reduction in dose to the left anterior descending coronary artery (LAD; shown in white in the figure), which is especially sensitive to radiation and a major factor in the development of radiation-induced heart disease in women who receive radiotherapy for cancer of the left breast.^{15,16} For this particular patient, the mean heart dose could be reduced from 3.0 Gy to 0.6 Gy, the LAD dose from 4.0 Gy to 0.6 Gy, and the lung dose from a mean value of 10.0 Gy to 6.5 Gy. As a result, the patient would be 20% less likely to develop posttreatment heart disease and 30% less likely to develop a secondary lung cancer.^{16,17} Similar potential improvements have been reported for other common cancers, such as brain tumors and prostate cancer. As a result, radiation oncologists' interest in *in vivo* imaging systems is growing rapidly, which we hope will help to spur their development and integration into routine clinical use.

Conclusions

In vivo imaging and range verification for proton- and ion-beam therapy are still in the R&D stage, but they could soon work their way into clinical use. Systems for in-beam PET are still undergoing initial testing and clinical trials; devices for prompt gamma imaging are just beginning to move into clinical testing; and research into ionoacoustic systems is ramping up in the wake of recent clinical

proof-of-principle studies. The proposed techniques visualize physical processes occurring on different time scales, and the performance of each is expected to vary based on anatomical location. Hybrid systems capable of measuring and imaging a combination of processes may therefore prove most beneficial for verifying treatment delivery. By reducing the uncertainty in beam delivery by millimeters or more, *in vivo* imaging could have a profound effect on doctors' ability to treat tumors, prevent life-altering posttreatment complications, increase cure rates, improve quality of life, and reduce health care costs incurred from management of posttreatment side effects for cancer patients.

References

1. R. Wilson, *Radiology* **47**, 487 (1946).
2. H. Paganetti, *Phys. Med. Biol.* **57**, R99 (2012).
3. W. Assman et al., *Med. Phys.* **42**, 567 (2015).
4. K. Parodi, W. Assman, *Mod. Phys. Lett. A* **30**, 1540025 (2015).
5. K. Parodi, *Nucl. Med. Rev. Suppl. C* **15**, C37 (2012).
6. J. F. Sutcliffe, *Phys. Med. Biol.* **41**, 791 (1996).
7. G. Sportelli et al., *Phys. Med. Biol.* **59**, 43 (2014).
8. T. Yamaya et al., *Phys. Med. Biol.* **56**, 1123 (2011).
9. C.-H. Min et al., *Appl. Phys. Lett.* **89**, 183517 (2006).
10. E. Testa et al., *Nucl. Instrum. Methods Phys. Res. B* **267**, 993 (2009).
11. M. Testa et al., *Phys. Med. Biol.* **59**, 4181 (2014).
12. J. M. Verburg, J. Seco, *Phys. Med. Biol.* **59**, 7089 (2014).
13. J. C. Polf et al., *Phys. Med. Biol.* **58**, 5821 (2013).
14. Y. Hayakawa et al., *Radiat. Oncol. Invest.* **3**, 42 (1995).
15. C. R. Correa et al., *J. Clin. Oncol.* **25**, 3031 (2007).
16. S. C. Darby et al., *New Engl. J. Med.* **368**, 987 (2013).
17. T. Grantzau et al., *Radiother. Oncol.* **111**, 366 (2014). ■

**Ultra Sensitive,
Deeply Cooled**

CCD Cameras



- Custom and OEM Scientific CCD and CMOS cameras
- Cryogenic and thermoelectrically cooled cameras
- Large area sensors; 2k x 2k, 4k x 4k, up to 10k x 10k CCDs
- Mosaic cameras with multiple sensors available
- Low noise, high dynamic range imaging

Contact us today to discuss your Custom and OEM camera design and manufacture!