

Particle Accelerators as Medical Devices

Marko Mehle, Luka Kurnjek
Cosylab d.d.
Gerbi
eva ulica 64
1000 Ljubljana
luka.kurnjek@cosylab.com



ABSTRACT: Particle accelerators have many applications; one of these that is gaining popularity is their use in medicine for cancer therapy. Proton and Carbon Ion therapy machines are used to accelerate a beam of particles and to deliver it very accurately into the tumor area. This kind of treatment is called Proton, Ion or simply Particle Therapy, and it presents several advantages over classic radiotherapy.

The accelerator machines involved in particle therapy are fairly modest either in power and complexity; yet the distance between the lab and the hospital is great – especially in paper work. One of the tasks that need to be fulfilled is to demonstrate that the machine is safe and effective, as required by the standards and regulations. The purpose of this article is to present a starting point for understanding what it takes to make a Particle Therapy Machine that is compliant with regulations.

Keywords: Medical Devices, Particle Accelerators, Accelerator Machine, Particle Therapy Machine, Cancer therapy, Tumor therapy

Received: 12 September 2018, Revised 5 November 2018, Accepted 7 November 2018

©2019 DLINE. All rights reserved

1. Introduction

1.1. Medical Devices

A particle accelerator, though complicated by itself, it is not enough to conduct ion therapy on patients, and it cannot be considered a “medical device”. Why? Because an accelerator is a machine that generates raw beam, without means of controlling or quantifying dose and position, and from a medical perspective is useless and even dangerous. The system that we call a “Particle Therapy Machine” is composed of different subsystems; in general: one to generate the beam, one to carry and direct it to the treatment area (for example, beam line and gantry), one to measure and control its delivery (dose delivery/scanning), one to accurately position the patient (robot and X-ray imaging), one to control the operation and execution of the treatment (treatment control system), and various systems for control and safety. They altogether constitute a “medical device”, as such it will have to be approved prior to its placement on the market. It is clear, for example, that it does not make sense to market a power supply magnet or a “dose control system” alone as medical devices since they do not have an intended use in medical treatment. And even if its purpose as a component of a therapy machine is undeniable, as in the case of a dose delivery system, it is not possible to judge its safety or effectiveness without knowing the properties of the remaining components of the system.

Nonetheless, quite often manufacturers of components and parts for medical accelerators deliver products that are “certification ready”, meaning that the component meets all applicable standards and regulations and that all necessary documentation is provided to the manufacturer of the Particle Therapy Machine.

2. Safety and Effectiveness

There are two distinct and equally important aspects of medical devices that every medical device needs to fulfill:

1. It must be safe.
2. It must be effective.

It's easy to understand why a Particle Therapy system needs to be safe: it is of no use to have a machine that cures cancer, while it is very likely to injure a patient every once in a while. At the same time, it is important for it to be effective, i.e. it is capable of fulfilling its medical purpose. Again, a system that it is safe but it does not do what it is supposed to do is useless. Imagine a wheelchair (also a medical device, though a simple one) that is tightly screwed to the floor: it may be as safe as it gets, but it does not allow the patient to move around, as it is supposed to.

That is why the laws regulating the market of medical devices take care of these two aspects. These are addressed in different ways. The manufacturer must demonstrate that his device is safe and effective, as required by the country where the device is to be sold. To prove safety, a lot of documentation must be provided: extensive test reports, detailed design documents, compliance with applicable standards, and so forth. In contrast, effectiveness is proven either by doing clinical trials and investigations involving animals or humans, or by claiming that the device in question is similar to other devices already being used, and demonstrating it by pointing to scientific articles or other literature.

3. Medical Requirements

The regulatory requirements that a medical device or component needs to fulfill are different from country to country, but quite often they end up being quite similar, since it is common that countries around the world recognize the use of international standards in their essence as a valid (and often necessary) mechanism to comply with their (local) regulations. All components of a Particle Therapy Machine can be very different in operation and purpose. Some come into contact with the patient or are installed within the treatment area, some not. Some are mechanical and have moving parts, others can be only electrical control and processing units. Some use high power or voltage, others not. Some contain or consist of software, some are fully analog or discretely digital. And the list goes on. There are lots of different regulatory requirements whose applicability depends on all these factors. But there are general things that are required for all “medical components”. Regardless of the above mentioned technical factors, a Quality Management System (such as [1]) Must be established, a Risk Management (as [2]) process must be defined and executed, and the safety and effectiveness of the component must be proven.

4. Industrial Requirements

Do all the subsystems or components of a Particle Therapy Machine need to be designed and manufactured according to the regulatory requirements for medical devices, or is there any possible exceptions? This is an important question, since devices that follow these are much more expensive than ordinary industrial stuff. The rule is that some of the components or subsystems may not need to comply with the applicable medical standards, as soon as it is possible to demonstrate that their malfunction can never lead to an unacceptable risk, for the patient and for other people involved. Even in cases where one of the subsystems is inherently critical to the safety or the essential performance of the whole system, it is possible to mitigate the potential risk by implementing alternative safety measures. The logic behind this approach is that we cannot really trust industrial (“non-medical”) devices since they were not designed and manufactured following quality processes required by medical standards, therefore it is assumed that they can contain bugs or can fail at any time. The mechanism to assess the initial risk and the suitability of the chosen risk mitigation measures is the Risk Management, usually at the system level.

A common example from the Particle Therapy industry is to build the beam generation subsystem (the particle accelerator itself) as an industrial device, and then to include in the system safety components (built as medical devices) that constantly monitor the beam parameters such as energy, current and position, and shutoff the beam if any of these is out of tolerance. This strategy may

help to save resources, but not always, since some complex industrial subsystems require safety measures sometimes so complex, that it is necessary to build a whole redundant instance of the original system, and this only to monitor the behavior of the original component. As an example, let's imagine an ion therapy system in which, besides the particle accelerator, the system that measures and controls the delivery of the dose to the patient is also a nonmedical device. In that case, the only possible way to mitigate the risk of patient overdose due to a failure would be to have an additional medical component that measures and monitors the delivered beam. Such redundancy would be unnecessary.

That is why we recommend writing the medical compliance strategy very early in the design process. (For those of you more interested in this aspect, I recommend taking a look at the standard IEC 60601- 1, section 16, and the informative Annex I).

5. Medical Software and Verification & Validation

Today, every relatively complex medical device contains at least some software. We all know that software is special, very different in nature and behavior to hardware. Even digital hardware systems composed of discrete logic are not comparable to software. The two main reasons for this are the complexity of software, with its enormous number of possible internal states, and the usual software development workflow itself, which is in practice more organic and less constrained than hardware design. The first factor determines that it is practically impossible to test a software module in all its internal states, leaving a room for latent bugs. The second factor also influences reliability, because since the nature of software allows for quick changes or last-minute fixes (that can go unnoticed), there is always a temptation to do it without the necessary care and to inadvertently introduce bugs.

If these two aspects are not addressed properly, software cannot be considered reliable enough for medical devices, especially when there is a risk to people associated with it. That's way the regulations and standards for medical devices take special care of software: to ensure that, although the possibility of bugs cannot be totally eliminated, their probability can be reduced and its associated risk mitigated; with good processes, proper testing, risk management and validation. There is an international standard that deals with this: "IEC62304: Medical device software — Software life cycle processes."

But software cannot exist without the associated hardware, and its proper operation in the context of the whole machine cannot be judged without testing them together. Yet, in a system that is as complex as a Particle Therapy Machine, there are many different subsystems or components that may contain software, and it is impractical or impossible to execute the full set of tests for every software module, all integrated and operating in the final machine. That is why this is never done in that way, and there is always a strategy to hierarchize and segregate testing into different levels, according to the system level architecture and the "V-model" (Figure 1), in order to make testing effective and practicable. This is the purpose of the set of activities usually referred to as "Verification and Validation": to *Verify*, at different system levels, that the build system works as specified, and to *Validate* that it is capable of fulfilling its intended purpose. More about verification and validation of Programmable Electrical Medical Systems can be found in the standard IEC60601-1, section 14, and informative annex H.




Figure 1. V-model

6. Purpose of Processes and Standards

Sometimes I hear, mostly amongst engineers and physicists, arguments about whether following the standards, regulations and processes can really enhance safety and effectiveness, or is it just a "placebo" to create an appearance of safety, which somehow

relieves the engineering conscience from the responsibility of thinking. The argument is valid and it has to be taken seriously. I find that for some people, for example, passing a list of tests automatically means that the device has no bugs, or that the relentless application of risk management can give you a 0-1 type of indication about the degree of safety of a device. As a consequence of the previous misconceptions, there are people that do what is called “safety by paperwork”, which means tweaking the documentation and doing ‘assessments’ until it is proven that the device is safe enough to meet the regulations. It is always possible to fool the system, and it is not in the scope of quality systems, processes and technical requirements to completely avoid it. Yet, what standards and regulations do over anything else is to establish clear responsibilities and minimum requirements, so – when things go wrong – no one can say “I didn’t know”.

On the other hand, it is important to understand that a good process helps you to keep organized through the development lifecycle. When projects are late or about to go over budget, there is a natural temptation to rush and skip steps. This may appear to save some time, but can also lead to dangerous mistakes. Having a well-established process naturally helps to resist the temptation. Furthermore, processes define what steps to take if there is an unexpected problem. And if the processes are wisely designed, they not only enhance safety but even make the development more efficient by keeping pace, consistency and making sure that no aspect or requirement is left neglected.

7. Conclusion

Building a Particle Therapy system is a technical challenge by itself. To make it compliant with standards and regulations it is yet another challenge, which is sometimes underestimated. Our experience shows that it is best to start designing the medical compliance strategy together with the machine itself and to already have solid processes established when the development and implementation start. Smart processes not only help to design a safe device and get through compliance; they may also increase productivity and decrease the uncertainties and defects, which may be, at the end of the day, the differentiators for staying in the game.

References

- [1] ISO 13485:2003: Medical devices — Quality management systems — Requirements for regulatory purposes (<https://www.iso.org/standard/36786.html>)
- [2] ISO 14971:2000: Medical devices — Application of risk management to medical devices (<https://www.iso.org/standard/31550.html>)