



(51) International Patent Classification:

A61N 5/10 (2006.01)

(21) International Application Number:

PCT/US2017/041076

(22) International Filing Date:

07 July 2017 (07.07.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/360,227 08 July 2016 (08.07.2016) US

(71) Applicant: MEVION MEDICAL SYSTEMS, INC.

[US/US]; 300 Foster St., Littleton, Massachusetts 01460 (US).

(72) Inventors: COOLEY, James; c/o Mevion Medical Sys-

tems, Inc., 300 Foster St., Littleton, Massachusetts 01460

(US). ZWART, Gerrit Townsend; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US). ROSENTHAL, Stanley J.; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US). JONES, Mark R.; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US). BOUCHET, Lionel G.; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US). CESSAC, Robert; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US). JACHINOWSKI, Joseph; c/o Mevion Medical Systems, Inc., 300 Foster St., Littleton, Massachusetts 01460 (US).

(74) Agent: PYSHER, Paul A. et al.; Choate, Hall & Stewart LLP, Two International Place, Boston, Massachusetts 02110 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,

(54) Title: TREATMENT PLANNING

Seven leaves on each side
in this example-




FIG. 30

(57) Abstract: An example method includes: receiving, from a treatment planning process, information that is based on a dose distribution for an irradiation target; and performing at least one of the following operations: moving structures to trim spots of a particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received; moving structures to produce a trimming curve for a layer of an irradiation target based on a specification of a trimming curve for the layer included in the information received; moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included in the information received; or moving structures based on configuration information for the structures in the information received.



CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

TREATMENT PLANNING

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of priority of U.S. Provisional Application No. 62/360,227, which was filed on July 8, 2016, and which is entitled "Treatment Planning". U.S. Provisional Application No. 62/360,227 is incorporated herein by reference.

TECHNICAL FIELD

[0002] This disclosure relates generally to treatment planning in a particle therapy system.

BACKGROUND

[0003] Particle therapy systems use an accelerator to generate a particle beam for treating afflictions, such as tumors. In operation, particles are accelerated in orbits inside a cavity in the presence of a magnetic field, and are removed from the cavity through an extraction channel. A magnetic field regenerator generates a magnetic field bump near the outside of the cavity to distort the pitch and angle of some orbits so that they precess towards, and eventually into, the extraction channel. A beam, comprised of the particles, exits the extraction channel.

[0004] The particle beam can damage healthy tissue adjacent to the irradiation target. A structure may be used to limit exposure of the healthy tissue to the particle beam. For example, the structure, or a portion thereof, may be placed in between the particle beam and the healthy tissue, thereby preventing exposure of the healthy tissue to the particle beam.

SUMMARY

[0005] An example method includes receiving, from a treatment planning process, information that is based on a dose distribution for an irradiation target; and configuring an adaptive aperture based on the information. The example method may include one or more of the following features, either alone or in combination.

[0006] The information received from the treatment planning process may identify structures of the adaptive aperture; and configuring the adaptive aperture may comprise moving the structures of the adaptive aperture to produce a trimming curve. The information may identify characteristics of a trimming curve; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce the trimming curve. The information may identify characteristics of pre-trimmed spots, which characteristics are stored in a library in computer memory; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to trim spots of a particle beam so that the spots of the particle beam approximate the pre-trimmed spots. The information may indicate where spots are to be formed on the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture based on where the spots are to be formed on the irradiation target. The information may identify characteristics of spots; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture on a spot-by-spot basis.

[0007] Configuring the adaptive aperture may comprise moving structures of the adaptive aperture so that at least two successive spots have at least one of a different size or a different shape. The information may comprise specifications of a trimming curve for a radiation field associated with the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce the trimming curve. The specifications may identify at least one of shape or a location of the trimming curve relative to a known point. The information may comprise specifications of distinct trimming curves for different layers of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce a trimming curve for each layer of the irradiation target based on specifications for the each layer.

[0008] The information received from the treatment planning process may comprise specifications of a single trimming curve for a radiation field associated with the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce the single trimming curve. The single trimming curve may be used to trim a radiation field for multiple layers of

the irradiation target. The information may comprise specifications for a configuration of the adaptive aperture for a radiation field of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture based on the specifications for the configuration of the adaptive aperture. The information may comprise specifications of configurations of the adaptive aperture for different layers of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture for the different layers of the irradiation target based on the information.

[0009] The adaptive aperture may comprise: structures that are movable relative to the irradiation target to produce a shape to trim part of a treatment area; and a primary carriage coupled to the structures to move the structures relative to the irradiation target in a first dimension relative to the irradiation target. Configuring the adaptive aperture may comprise moving at least one of a structure among the structures or the primary carriage. The adaptive aperture may comprise secondary carriages coupled to the primary carriage and to the structures to move the structures in a second dimension relative to the irradiation target. Configuring the adaptive aperture may comprise moving at least one of the secondary carriages. The second dimension may comprise movement into, and out of, the treatment area. Moving at least one of the secondary carriages may comprise rotating the at least one secondary carriage. The structures may comprise leaves that are extendible and retractable relative to the particle beam. Configuring the adaptive aperture may comprise rotating all or part of the adaptive aperture.

[0010] An example method comprises receiving, from a treatment planning process, information that is based on a dose distribution for an irradiation target; and performing at least one of the following operations: moving structures to trim spots of a particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received; moving structures to produce a trimming curve for a layer of an irradiation target based on a specification of a trimming curve for the layer included in the information received; moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included

in the information received; or moving structures based on configuration information for the structures in the information received. The structures may be part of an adaptive aperture and may comprise leaves that are movable into, or out of, an area to be treated by the particle beam. The information may specify a configuration of the leaves and/or carriage(s) holding the leaves.

[0011] An example particle therapy system comprises: a particle accelerator to output a particle beam for application to an irradiation target; an adaptive aperture between the particle accelerator and the irradiation target, where the adaptive aperture comprises structures that are movable into, and out of, a path of the particle beam; and a control system comprising one or more processing devices configured to perform various operations. The various operations may comprise: receiving, from a treatment planning process, information that is based on a dose distribution for the irradiation target; and configuring the adaptive aperture based on the information. The example particle therapy system may include one or more of the following features, either alone or in combination.

[0012] The information based on the dose distribution may identify characteristics of pre-trimmed spots, with the characteristics being stored in a library in computer memory; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to trim spots of a particle beam so that the spots of the particle beam approximate the pre-trimmed spots. The information may indicate where the pre-trimmed spots are to be formed on the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture based on where the pre-trimmed spots are to be formed on the irradiation target. The characteristics may comprise one or more parameters, with the one or more parameters comprising at least one of a radius associated with a spot or an angle of rotation of the spot. The information may identify characteristics of spots; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture on a spot-by-spot basis.

[0013] Configuring the adaptive aperture may comprise moving structures of the adaptive aperture so that at least two successive spots have at least one of a different size or a different shape. Configuring the adaptive aperture may comprise

maintaining structures of the adaptive aperture so that at least two successive spots have at least one of a same size or a same shape. The information may comprise specifications of a trimming curve for a radiation field associated with the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce the trimming curve. The specifications may identify at least one of shape or a location of the trimming curve relative to a known point.

[0014] The information based on the dose distribution may comprise specifications of trimming curves for different layers of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce a trimming curve for each layer of the irradiation target based on specifications for the each layer. The information may comprise specifications of a single trimming curve for a radiation field associated with the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture to produce the single trimming curve, with the single trimming curve being used to trim a radiation field for multiple layers of the irradiation target.

[0015] The information may comprise specifications for a configuration of the adaptive aperture for a radiation field of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture based on the specifications for the configuration of the adaptive aperture. The information may comprise specifications of configurations of the adaptive aperture for different layers of the irradiation target; and configuring the adaptive aperture may comprise moving structures of the adaptive aperture for the different layers of the irradiation target based on the information. The adaptive aperture may comprise a primary carriage coupled to the structures to move the structures relative to the irradiation target in a first dimension relative to the irradiation target; and configuring an adaptive aperture may comprise moving at least one of a structure among the structures or the primary carriage. The adaptive aperture may comprise secondary carriages coupled to the primary carriage and to the structures to move the structures in a second dimension relative to the irradiation target; and configuring the adaptive aperture may comprise moving at least one of the secondary carriages.

[0016] The second dimension may comprise movement into, and out of, the treatment area. Moving at least one of the secondary carriages may comprise rotating the at least one secondary carriage. The structures may comprise leaves that are extendible and retractable relative to the particle beam. Configuring the adaptive aperture may comprise rotating all or part of the adaptive aperture.

[0017] An example particle therapy system comprises: a particle accelerator to output a particle beam for application to an irradiation target; and a control system comprising one or more processing devices configured to perform various operations. The operations comprise receiving, from a treatment planning process, information that is based on a dose distribution for the irradiation target; and performing at least one of the following operations: moving structures to trim spots of the particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received; moving structures to produce a trimming curve for a layer of an irradiation target based on a specification of a trimming curve for the layer included in the information received; moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included in the information received; or moving structures based on configuration information for the structures in the information received. The structures may be part of an adaptive aperture and may comprise leaves that are movable into, or out of, an area to be treated by the particle beam.

[0018] An example method that may be performed by a particle therapy system comprises obtaining a target dose distribution for a treatment area; identifying one or more pre-trimmed spots from among a library of pre-trimmed spots that best approximates one or more points of the target dose distribution; and outputting information identifying the one or more pre-trimmed spots. The example method may comprise one or more of the following features either alone or in combination.

[0019] The example method may comprise obtaining a set of coarse spots that approximates the target dose distribution; changing an intensity of one or more spots in the coarse set based on a maximum acceptable dosage for a portion of the

treatment area; identifying a spot among the coarse spots that is within a transition region of the treatment area; and substituting the one or more pre-trimmed spots for the spot within the transition region. The example method may comprise obtaining a minimum acceptable dosage for a portion of the treatment area; where the set of coarse spots is selected based on the minimum acceptable dosage. The information may be usable to configure an adaptive aperture to achieve a dose distribution during treatment that approximates the target dose distribution.

[0020] An example method that may be performed by a particle therapy system comprises obtaining treatment planning information; determining that an irradiation target has moved since a prior application of radiation based on the treatment planning information; modifying the treatment planning information based on movement of the irradiation target to produce updated treatment planning information; and applying radiation to the irradiation target based on the updated treatment planning information. The example method may comprise one or more of the following features either alone or in combination.

[0021] Modifying the treatment planning information may be performed by a computing system of the particle therapy system. Determining that that irradiation target has moved may be based on a comparison of images of the irradiation target captured between two points in time. The treatment planning information may be modified by changing locations to which spots of radiation are to be applied to the irradiation target based on a location of the irradiation target following movement.

[0022] Two or more of the features described in this disclosure, including those described in this summary section, may be combined to form implementations not specifically described herein.

[0023] Control of the various systems described herein, or portions thereof, may be implemented via a computer program product that includes instructions that are stored on one or more non-transitory machine-readable storage media, and that are executable on one or more processing devices (e.g., microprocessor(s), application-specific integrated circuit(s), programmed logic such as field programmable gate array(s), or the like). The systems described herein, or portions thereof, may be implemented as an apparatus, method, or electronic system that

may include one or more processing devices and computer memory to store executable instructions to implement control of the stated functions.

[0024] The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF THE DRAWINGS

[0025] Fig. 1 is a perspective view of an example adaptive aperture leaf that is usable with the example adaptive apertures described herein.

[0026] Fig. 2 is a top view of an example adaptive aperture leaves positioned relative to a treatment area of an irradiation target.

[0027] Fig. 3 is a perspective view of an example adaptive aperture.

[0028] Fig. 4 is a front view of an example adaptive aperture having components portrayed in see-through to show the interiors thereof.

[0029] Fig. 5 is a perspective view of an example adaptive aperture having components portrayed in see-through to show the interiors thereof.

[0030] Fig. 6 is a top view of example control structures used to move a leaf into, and out of, a treatment area of an irradiation target.

[0031] Fig. 7 is a perspective view of an example adaptive aperture having curved leaves.

[0032] Fig. 8 is a cross-sectional view of an example synchrocyclotron for use in a particle therapy system.

[0033] Fig. 9 is a cross-sectional view of an example synchrocyclotron for use in a particle therapy system.

[0034] Fig. 10 is a side view of an example scanning system.

[0035] Fig. 11 is a perspective view of components of an example scanning system.

[0036] Fig. 12 is a front view of an example magnet for use in a scanning system of the type shown in Figs. 10 and 11.

[0037] Fig. 13 is a perspective view of an example magnet for use in a scanning system of the type shown in Figs. 10 and 11.

[0038] Fig. 14 is a perspective view of an example energy degrader (range modulator) for use in a scanning system of the type shown in Figs. 10 and 11.

[0039] Fig. 15 is a perspective view of a process for moving a leaf of an energy degrader in the path of a particle beam.

[0040] Fig. 16 is a perspective view of an adaptive aperture positioned relative to a patient during particle therapy treatment.

[0041] Fig. 17 is a flowchart showing an example process for performing raster scanning that may be performed using the hardware of Figs. 10 to 15.

[0042] Fig. 18 is a top view show an example cross-section of an irradiation target and a radiation scan path.

[0043] Fig. 19 is a perspective view of an example therapy system.

[0044] Fig. 20 is an exploded perspective view of components of an example synchrocyclotron for use in the particle therapy system.

[0045] Fig. 21 is a cross-sectional, top view of the example synchrocyclotron.

[0046] Fig. 22 is a perspective view of the example synchrocyclotron.

[0047] Fig. 23 is a cross-sectional view of an example ion source for use in the synchrocyclotron.

[0048] Fig. 24 is a perspective view of an example dee leaf and an example dummy dee for use in the synchrocyclotron.

[0049] Fig. 25 is a conceptual view of an example particle therapy system that may use a variable-energy particle accelerator.

[0050] Fig. 26 is a perspective, exploded view of an example magnet system that may be used in a variable-energy particle accelerator.

[0051] Fig. 27 is an example graph showing energy and current for variations in magnetic field and distance in a particle accelerator.

[0052] Fig. 28 is a side view of an example structure for sweeping voltage on a dee leaf over a frequency range for each energy level of a particle beam, and for varying the frequency range when the particle beam energy is varied.

[0053] Fig. 29 is a top view of example adaptive aperture leaves on carriages mounted on a track that is curved relative to the direction of a particle beam.

[0054] Fig. 30 is a top view of example adaptive aperture leaves positioned relative to a treatment area of an irradiation target.

[0055] Fig. 31 is a graph showing an Gaussian distribution of a spot produced by the example particle therapy system.

[0056] Fig. 32 depicts an original spot and a pre-trimmed spot produced by trimming the original spot.

[0057] Fig. 33 is a perspective view of an irradiation target containing layers to be treated by particle therapy.

[0058] Fig. 34 is a perspective view of an irradiation target containing layers to be treated, and application of a trimming curve to the irradiation target.

[0059] Fig. 35 is perspective view of a layer of an irradiation target and trimming curves superimposed thereon.

[0060] Fig. 36 is a flowchart showing an example scanning process that includes configuring trimming curves based on information about a target dose distribution.

[0061] Fig. 37 is a flowchart showing an example of an optimization process performed to model, and to obtain, treatment plan information.

[0062] Fig. 38 is a side view of an example treatment area that identifies dosage to be applied to the treatment area.

[0063] Fig. 39 is a side view of a coarse set of spots to be applied to the treatment area.

[0064] Fig. 40 is a side view of the set of spots of Fig. 39, with a spot having been intensity modulated.

[0065] Fig. 41 is a side view of the set of spots of Fig. 40, with a spot having been replaced by a pre-trimmed spot.

[0066] Fig. 42 shows side views of example spots in a library.

[0067] Fig. 43 is a flowchart showing an example process for configuring an adaptive aperture in response to movement of an irradiation target.

[0068] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0069] Described herein are example implementations of an adaptive aperture (also called a “configurable collimator”) that may be used to control the extent of radiation, such as a proton or ion beam, applied to a patient. In some implementations, an adaptive aperture includes structures that are controllable to allow some radiation to pass to a patient and to block some radiation from passing to the patient. Typically, the radiation that passes is directed to an irradiation target to be treated, and the radiation that is blocked would otherwise hit, and potentially damage, healthy tissue. Radiation internal to the target may be blocked as well in order to achieve a target dose distribution. In an example operation, the adaptive aperture is placed in the radiation path between a radiation source and an irradiation target and is controlled to produce a trimming curve of an appropriate size and shape. A trimming curve may be an open-ended curve or a closed curve, such as an opening that allows some radiation to pass to the irradiation target, while preventing some radiation from reaching adjacent tissue. The adaptive aperture may be used in any appropriate radiation therapy system, and is not limited to use with any particular type of system, including those described herein. Examples of systems with which the adaptive aperture may be used are described herein.

[0070] In some implementations, the adaptive aperture contains generally flat structures, which are referred to as “plates” or “leaves”, and which are controllable to move into the “beam” or “treatment” area to block some radiation and allow some radiation to pass. The leaves are controllable to produce an opening of size and shape that is appropriate for the current treatment. In some implementations, the leaves are held on two carriages, and face each other and the treatment area. The leaves are controllable to move into, and out of, the treatment area, to form the opening (or aperture) through which the particle beam passes to treat an area of a patient. The leaves that form the opening also block radiation from passing to tissue (e.g., healthy tissue) adjacent to the opening that is covered by the leaves. In this context, covering includes the leaves blocking the particle beam. The leaves, or a portion thereof, may be placed in between the particle beam and target tissue to

create sharp dose gradients which may better optimize the dose distribution to satisfy an objective function when combined with other particle beams.

[0071] In some implementations, each set of leaves is configurable to define an edge that is movable into a path of the particle beam so that a first part of the particle beam on a first side of the edge is blocked by the leaves, and so that a second part of the particle beam on a second side of the edge is not blocked by the leaves and is allowed to pass to the treatment area. In some implementations the leaves are connected to, are part of, or include, motors – one per leaf – that are controllable to control movement of the leaves towards or away from the treatment area to define the edge. In some implementation, the motors may be linear motors. An example linear motor includes a movable component and a stationary component. The stationary component may include a magnetic field generator, such as two magnets having their poles aligned, to generate a first magnetic field. The movable component may be sandwiched between the two magnets and may include one or more coils to conduct current to produce a second magnetic field that interacts with the first magnetic field to cause the moveable component to move relative to the stationary component. The movable component of each linear motor is connected to, or part of, a corresponding one of the leaves such that the corresponding leaf moves along with movement with the movable component.

[0072] Fig. 1 shows an example of a leaf 40 that may be used in the adaptive aperture, although the adaptive aperture is not limited to use with this type of leaf. The height 50 of the leaf is along the beamline (e.g., the direction of the particle beam). The length 52 of the leaf is along its direction of actuation into and out of the treatment area, and is based on the field size, or portion thereof, that the system can treat. The field size corresponds to the treatment area that the beam can impact. The width 53 of the leaf is the direction along which multiple leaves stack when actuated. Generally, the more leaves that are used, the higher the resolution of the trimming curve that can be produced, including for curved boundaries.

[0073] In Fig. 1, leaf 40 includes a tongue and groove feature 55 along its side, which is configured to reduce inter-leaf leakage when multiple such leaves stack. In this example, the curved end 56 of leaf 40 is configured to maintain a

surface tangent to the beam at all locations in the treatment area. However, in some implementations, the end of each leaf may be flat, not curved.

[0074] In some implementations, the adaptive aperture leaves have a height that is sufficient to block at least the maximum beam energy (e.g., the maximum energy of the particle beam output by a particle therapy system). In some implementations, the adaptive aperture leaves have a height that blocks less than the maximum beam energy, for reasons described below. In some implementations, the adaptive aperture leaves have lengths that are dictated not by the area of an entire treatment area, but rather by the area of a single beam spot or multiple beam spots. In this context, a “beam spot” includes the cross-sectional area of a particle beam. In some implementations, the adaptive aperture is controllable to move into the “beam” or “treatment” area to block some radiation and allow some radiation to pass from an individual or multiple beam spot(s).

[0075] In an example, the particle therapy system may be configured to treat a tumor having a cross-section that can fit into a 20cm x 20cm square area. In this example, each leaf in the adaptive aperture may have a length of about 2cm, which is about enough to block particles in half of one beam spot. As noted, the adaptive aperture may include sets of leaves that face each other. So, leaves from each set may be controlled to cover the whole, single beam spot, if necessary, thereby preventing passage of radiation. The leaves may also be controllable to create one or more trimming curves that define an opening through which some, or all, of the radiation from the single beam spot can pass.

[0076] In operation, the adaptive aperture is configured to move as the beam scans across the radiation target, and to track the beam’s movement during scanning. In an example, the adaptive aperture may be configured to move about 20cm so as to enable coverage over the entirety of the 20cm x 20cm area. As described above, the adaptive aperture may be configured to use enough leaves to cover (or “trim”) one beam spot and, in some case, a small amount of extra area (e.g., 5% extra area, 10% extra area, 15% extra area, or 20% extra area).

[0077] Fig. 2 shows an example implementation of an adaptive aperture 700. In this example, adaptive aperture 700 includes leaves 701 having a height and

made of a material, such as nickel, brass, tungsten, or other metal, sufficient to inhibit or prevent passage of radiation at a given energy. For example in some systems, a particle accelerator is configured to generate a particle beam having a maximum energy of 100 MeV to 300 MeV. Accordingly, in such systems, the leaves may be constructed to prevent passage of a beam having an energy of 100 MeV, 150 MeV, 200 MeV, 250 MeV, 300 MeV, and so forth.

[0078] Leaves 701 are mounted on carriages to control their movement relative to a treatment area of an irradiation target, such as a cross-sectional layer of a tumor in a patient. The movement is controlled to cause leaves 701 to cover some parts of treatment area 704, thereby preventing radiation from impacting those parts during treatment, while leaving other parts of treatment area exposed to the radiation. In the example implementation of Fig. 2, there are fourteen leaves in total – seven on the left and seven on the right. In some implementations, there may be a different number of leaves, e.g., ten in total, five on the left and five on the right, twelve in total, six on the left and six on the right, and so forth.

[0079] The adaptive aperture may be used with any appropriate type of radiation therapy system. In an example implementation, the radiation therapy system is a proton therapy system. As described herein, an example proton therapy system scans a proton beam across a treatment area of an irradiation target in order to destroy malignant tissue. During scanning, the particle beam moves across the treatment area to cover the treatment area with radiation. In an example implementation, the particle beam is pulsed. Even with pulsed beams, uniform dose coverage can be achieved with a single pass of the beam. Each spot may be hit with multiple pulses to account for random variability in the intensity of each pulse that cannot be controlled. The particle beam may be passed over the treatment area multiple times in order to mitigate organ motion effects due to patient motion (e.g., due to breathing). As a result, it may be beneficial to scan the same treatment area more than once. Each successive scan may be offset from one or more other scan(s). An example of this type of scanning is called pencil-beam scanning, and the repetitive scans are referred to as painting or repainting the treatment area.

[0080] The irradiation target is typically three-dimensional in structure. Accordingly, as described herein, the irradiation target is treated cross-sectional layer (or simply “layer”) by layer. That is, a layer of the irradiation target is treated, followed by another treatment of another layer, and so forth until the entire target is treated. Different layers of an irradiation target are treated by varying an energy level of the particle beam. That is, different energy levels of the particle beam impact different layers of the irradiation target, with higher energy levels affecting layers deeper inside the irradiation target relative to the particle beam source. Accordingly, during treatment, the energy level of the particle beam is changed in order to reach, and thus treat, different layers of the irradiation target.

[0081] Fig. 2 shows leaves 701 configured to allow radiation to impact part of a layer (e.g., the treatment area) and to prevent radiation from impacting other parts of the layer (e.g., healthy tissue). In Fig. 2, locations 707 represent centers of beam spots to be delivered during a scan of a proton beam across treatment area 704. Circle 708 represents a treatment boundary beyond which no radiation is intended to be delivered. Beam spots that are close to this boundary (e.g., within one standard deviation of the particle beam’s profile) border healthy tissue. These are spots that are trimmed (that is, blocked) by appropriate configuration and placement of leaves on the adaptive aperture. An example of a beam spot to be trimmed is beam spot 711, having its center at location 706. As shown, one of leaves 701 is configured to block the portion of beam spot 711 that extends beyond circle 708.

[0082] In an example implementation, on each of two separate carriages, there are five leaves that are about 5mm in width and two leaves that are about 20mm in width, for a total of seven leaves per carriage. Thus, in some implementations, on each of two separate carriages, there are seven leaves, two of which each have widths that are three times or more the widths of each of five other leaves. Other implementations may contain different numbers, sizes, and configurations of leaves, and different numbers and configurations of carriages. For example, some implementations may include any number between five and fifty leaves per carriage, e.g., 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21,

22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, or 50 leaves (or more) per carriage.

[0083] The carriages can move both horizontally and vertically, as described herein. The leaves are also movable horizontally relative to each carriage into, and out of, the treatment area. In this way, the leaves are configurable to approximate the shape of the treatment boundary in the region near the area being treated (e.g., circle 711 or a portion thereof in this example).

[0084] The leaves may be moved vertically and/or horizontally between different scans of the particle beam so that the leaves are in appropriate positions when the beam is delivered to a particular area. The leaves need not necessarily be moved for every scan pass, but instead may be moved to a location appropriate for an area. In some cases, e.g., for spots interior to the treatment area, radiation treatment may proceed without the trimming provided by the adaptive aperture.

[0085] Fig. 30 shows another example of leaves 1400 that are part of an adaptive aperture configured to trim a radiation spot 1401 centered at location 1402. In this example, there are seven leaves on each of two sides of the adaptive aperture (supported by corresponding carriages). The leaves on each side include two leaves that are wider than the other five; however, the adaptive aperture is not limited to this configuration. In this case, spot 1401 has a radius of 2.5 sigma defining an 8mm Gaussian radiation spot.

[0086] Figs. 3, 4, and 5 show an example implementation of components of an example implementation of an adaptive aperture, including carriages 713, 714, 715 configured to hold, and to move, the leaves described above both vertically and horizontally relative to the treatment target. As shown, vertical movement includes movement in the Cartesian Z-dimension 717, and horizontal movement includes movement in the Cartesian X-dimension 718 (with the Cartesian Y-dimension being into, or out of, the page in Fig. 4). Figs. 4 and 5 show parts of carriage housings as transparent in order to show components inside the housings; however, the housings are not actually transparent.

[0087] Carriage 713 is referred to herein as the primary carriage, and carriages 714 and 715 are referred to herein as secondary carriages. Secondary

carriages 714, 715 are coupled to primary carriage 713, as shown in Figs. 3 to 5. In this example, secondary carriages 714, 715 each include a housing that is fixed to primary carriage 713 via a corresponding member 718, 719. In this example, primary carriage 713 is movable vertically (the Z-dimension) along tracks 720. The vertical movement of primary carriage 713 also causes the secondary carriages to move vertically. In some implementations, the secondary carriages move vertically in concert. In some implementations, vertical movement of each secondary carriage is independent of vertical movement of the other secondary carriage.

[0088] As shown in Figs. 3 to 5, each secondary carriage 714, 715 is connected to a corresponding rod 722, 723, along which the secondary carriage moves. More specifically, in this example, motor 725 drives secondary carriage 714 to move along rod 722 towards or away from secondary carriage 715. Likewise, in this example, motor 726 drives secondary carriage 715 to move along rod 723 towards or away from secondary carriage 714. Control over movement of the primary and secondary carriages is implemented to position the leaves relative to the irradiation target, as described herein. In addition, the leaves themselves are also configured to move in and out of the carriages, as also described herein.

[0089] As shown in Fig. 5, a motor 730 drives the vertical movement of primary carriage 713. For example, as shown in Fig. 4, lead screw 731 is coupled to housing 732, which holds motors 725, 726 that drive corresponding secondary carriages 714, 715, and which is mounted on tracks 720. Lead screw 731 is coupled to, and driven in the Z-dimension by, motor 730. That is, motor 730 drives lead screw 731 in the Z-dimension. Because lead screw 731 is fixed to housing 732, this movement also causes housing 732, and thus secondary carriages 714, 715, to move along tracks 720 in the Z-dimension.

[0090] In this example implementation, as noted, seven leaves 735, 736 are mounted on each secondary carriage 714, 715. Each secondary carriage may be configured to move its leaves horizontally (in the X-dimension) into, or out of, the treatment area. The individual leaves on each secondary carriage may be independently movable in the X-dimension relative to other leaves on the same secondary carriage. In some implementations, the leaves may also be configured to

move in the Y-dimension. Furthermore, the leaves on one secondary carriage 714 may be movable independently of the leaves on the other secondary carriage 715. These independent movements of leaves on the secondary carriages, together with the Z-dimension movements enabled by the primary carriage, allow the leaves to be moved into various configurations. As a result, the leaves can conform, both horizontally and vertically, to treatment areas that are arbitrarily shaped both in horizontal and vertical dimensions. The sizes and shapes of the leaves may be varied to create different conformations.

[0091] The leaves may be made of any appropriate material that prevents or inhibits transmission of radiation. The type of radiation used may dictate what material(s) are used in the leaves. For example, if the radiation is X-ray, the leaves may be made of lead. In the examples described herein, the radiation is a proton or ion beam. Accordingly, different types of metals or other materials may be used for the leaves. For example, the leaves may be made of nickel, tungsten, lead, brass, steel, iron, or any appropriate combinations thereof. The height of each leaf may determine how well that leaf inhibits transmission of radiation.

[0092] In some implementations, the leaves may have the same height, whereas in other implementations, some of the leaves may have heights that are different from heights of others of the leaves. For example, in Figs. 2 to 5, the leaves are each 5mm in height. However, any appropriate heights may be used. For example, leaves 735, 736 may have any of the following (or other heights): 1mm, 2mm, 3mm, 4mm, 5mm, 6mm, 7mm, 8mm, 9mm, 10mm, 11mm, 12mm, 13mm, 14mm, 15mm, 16mm, 17mm, 18mm, 19mm, 20mm, 21mm, 22mm, 23mm, 24mm, 25mm, 26mm, 27mm, 28mm, 29mm, and so forth. The leaves may have any combination of the foregoing heights. In addition, each of the leaves may have a different height than one or more others of the leaves.

[0093] In some implementations, leaves of the same width or of varying widths may be stacked to inhibit radiation transmission. In some implementations, thinner leaves (e.g., leaves of lesser width) may be used in connection with thicker leaves (e.g., leaves of greater width). In some implementations, the leaves have enough height to completely block the particle beam at the maximum beam energy.

Leaves of the same height or of varying heights may be stacked to inhibit radiation transmission. In some implementations, shorter leaves (e.g., leaves of lesser height) may be used in connection with longer leaves. In some implementations, the leaves have enough height to completely block the particle beam at the maximum beam energy. In some implementations, the leaves have enough height to block the particle beam at less than the maximum beam energy (and not to block the particle beam at maximum energy). For example, while a proton therapy system may have the capability to deliver a beam of energy 230 MeV that can treat to a depth in the patient of 32cm, in some implementations, the adaptive aperture can only block protons of at most 175 MeV that can treat to a depth of no more than 20cm. In doing so, less beam-stopping material may be used, e.g., 2.1cm of Tungsten instead of 3.3cm, or 3.3cm of Nickel instead of 5.2cm. In this example, the proton therapy system would still be capable of treating at depths exceeding 20cm, but the adaptive aperture would not be used for such treatments. This may be deemed acceptable because, in some circumstances, deeper treatments benefit less from the particle beam collimation that the adaptive aperture provides. That is, in some treatment scenarios, shallow, low-energy treatments are where the adaptive aperture will be most effective, and there may be engineering advantages to reducing the amount of material in the leaves. Thus, in some example implementations, shorter leaves are used and the adaptive aperture is restricted to use with shallow, lower-than-maximum energy treatments.

[0094] In the implementations of Figs. 2 to 5, the leaves are semi-rectangular in shape and have about the same surface area when viewed from the side. In some implementations, this need not be the case. For example, the leaves may have different shapes than those shown. Example shapes include, but are not limited to, circular shapes, curvilinear shapes, oval shapes, square shapes, and triangular shapes. Furthermore, individual leaves may have different shapes than others of the leaves contained in the same carriage or in a different carriage. For example, one carriage may contain both rectangular and curvilinear shaped leaves.

[0095] In some implementations, the leaves have heights that are enough not only to fully stop a particle beam at the maximum expected proton energy (e.g.,

3.3cm of Tungsten at 230 MeV or, e.g., 5.2cm of nickel), but also to have enough extra material to prevent proton transmission between the leaves. This material may have a tongue and groove structure as shown in Fig. 1, or a similar configuration. The leaf ends may be configured to include curved or tapered surfaces to enhance delivered penumbra for proton beams of various divergence.

[0096] In some implementations, there may be more than one primary carriage and corresponding motors and rails. For example, a first primary carriage may control Z-dimension movement of a first secondary carriage, and a second primary carriage may control Z-dimension movement of a second secondary carriage. Therefore, in such implementations, the two secondary carriages can be moved independently in the Z-dimension, if desired. In any case, the primary carriage may be computer controlled. For example, executable instructions are stored in computer memory (e.g., one or more non-transitory machine-readable storage media), and executed by one or more processing devices to control the movement. Control may be performed with, or without, user input during treatment.

[0097] As explained, each secondary carriage 714, 715 includes a corresponding motor to control X carriage movement, as described above. In some implementations, there may be also one motor for each leaf, whereas in other implementations a single motor may control all leaves. The motors may be mounted on a corresponding secondary carriage. As a result, the motors move along with the corresponding secondary carriage. As noted, the motors control movement of leaves in each carriage. The leaves may be mounted on actuators that enable movement in two dimensions. In the examples of Figs. 2 to 5, the leaves are each part of an integrated leaf. Referring to the example of Fig. 6, mechanisms enable the individual leaves to move in and out (the X-dimension) of the integrated leaf that comprises each carriage into, or out of, the treatment area.

[0098] In some implementations, all leaves on a single carriage are independently movable. Fig. 6 shows the movement mechanism for a leaf 735a that is part of a secondary carriage. Each leaf in each secondary carriage may have a configuration, and movement mechanism, that is similar to, or identical to, the leaf and movement mechanism of Fig. 6.

[0099] In the example of Fig. 6, the movement mechanism includes a leaf stick 735b, a lead screw nut 735c, a lead screw 735d, a thrust bearing assembly 735e, a coupling 735f, a motor 735g, a bearing block 735h, a sleeve bearing pin 735i, a sleeve bearing 735j, and a motor mount block 735k. In operation the motor shaft, through the coupling, spins the lead screw. This causes the lead screw nut to advance or retract depending on the direction of screw rotation. The lead screw nut is fixed to the leaf assembly so the leaf is moved forward or backward (into or out of the beam path) as the motor spins. Sleeve bearing pins are fixed to the leaf and slide along sleeve bearings which are held captive in the bearing block that supports the whole assembly. This bearing block houses the bearings and provides space for the motor assemblies for all leaves on the carriage.

[00100] In some implementations, the motors described with respect to Fig. 6 may be replaced with linear motors. In some implementations, the linear motors are controllable to configure a set of leaves to define a first edge, and to configure another set of leaves to define a second edge that faces the first edge. Each of the linear motors may include a movable component and a stationary component. The stationary component may include a magnetic field generator to generate a first magnetic field. An example of a magnetic field generator includes two stationary magnets that are adjacent and spaced apart, and that have their poles aligned. The movable component includes one or more coils to conduct current to produce a second magnetic field that interacts with the first magnetic field to cause the moveable component to move relative to the stationary component. For example, the movable component may be a coil-carrying plate between the two magnets that make up the stationary component. When current passes through the coil, that current produces a magnetic field that interacts with the magnetic field produced by the two magnets, and that causes the movable component (e.g., the current-carrying plate) to move relative to the two magnets. In some implementations, a single linear motor may control a single leaf. Because a leaf is attached to the movable component, the leaf moves along with the movable component. The linear motors of different leaves may be controlled to control movement of the leaves, and thus to define the edges of the adaptive aperture described herein.

[00101] As described, the adaptive aperture can be used to trim each layer of a scanned treatment with a different aperture shape, allowing for three-dimensional field shaping techniques such as layer stacking. However, an implementation of the adaptive aperture may be configured to mimic either a machined patient-specific structure or a multi-leaf collimator.

[00102] In this regard, some existing treatment planning systems (TPSs) perform treatment planning treatment planning processes that calculate the shape of a fixed brass aperture meant to be uniformly applied to the entire treatment volume. TPSs may also dictate, or contain instructions to configure, an adaptive aperture shape for a target. A computer program can interpret an aperture shape, either from a continuous aperture curve or set of fixed multi-leaf collimator leaf positions, and translate that shape into a series of dynamic leaf positions for the adaptive aperture that correlate with delivered radiation treatment spots. Existing TPS functionality can thus be leveraged, allowing the adaptive aperture to be compatible, with relatively few modifications, to existing TPS software.

[00103] Another example technique that may improve edge conformality and that is enabled by the adaptive aperture takes advantage of the possibility that a treatment dose to the same treatment area of a layer will be delivered in several passes or paintings. Repainting, which is described herein, is a technique used in scanned proton therapy in which dose is divided to a treatment area using several particle beam passes in order to spread the delivered radiation uniformly during a timescale that is long compared with patient motion, such as breathing. There are also other reasons in which dose to a given spot may not all be delivered in one pulse. Some example reasons include: dynamic dose control can regulate the charge in each pulse to deliver an accurate total dose, there are safety limits to the charge in a pulse, and a required dynamic range of dose may exceed the capabilities of the particle accelerator.

[00104] If the spots in an irradiated volume will be painted several times during a treatment (e.g., during multiple scans), then the conformality of the adaptive aperture to the lateral edge of the irradiation target can be improved by shifting the leaves slightly (e.g., at the sub-millimeter level), vertically, horizontally, or both

vertically and horizontally between each painting (e.g., between each scan pass of the particle beam covering a same treatment area). In this way, what might have been a slightly jagged edge owing to the finite size of the individual leaves can be smoothed to attain a better approximation of the intended aperture curve.

[00105] Adding rotational degrees of freedom can improve the ability of the adaptive aperture to conform to radiation targets. For example, the entirety of the assembly of Figs. 3 to 5 may be configured to be rotated in a plane perpendicular to the beam direction, in a plane parallel to the beam direction, or in a combination thereof. In some implementations, each individual secondary carriage 714, 715 may be configured to rotate independently in a same plane. In this way, the adaptive aperture may provide more flexibility to conform to complex shapes that are not ideally oriented. In some implementations, both the primary carriage and each secondary carriage may be rotatable.

[00106] In the example implementations described above, each leaf is independently actuated such that an arbitrary shape can be traced with a leaf configuration. It may be, however, that such flexibility is not required to achieve acceptable edge conformality. The leaves could be mechanically constrained with the ability to achieve only a finite number of configurations. For example, the leaves could be restricted to arrangements that put them in a vertical line, forward diagonal shape, backward diagonal shape, concave shape, convex shape, or any other achievable shape. In this way, flexibility could be traded for mechanical simplicity. The leaves may be actuated using a cam timing gear through one remote motor rather than four motors mounted on a secondary carriage. In some implementations, discrete shapes of the gears may be machined out of sectors of wheels which are rotated into place as needed. In order to reduce the number of motors, feedback, controllers, and associated wiring, an assembly may be used that has one motor that achieves discrete leaf configurations.

[00107] Fig. 7 shows an example implementation of an adaptive aperture having cams that are rotatable relative to shaft 751. In the example of Fig. 7, all leaves 750 are actuated by one motor. This motor is at the top of vertical shaft 751.

The shape and orientations of leaf cams are chosen to achieve different leaf configurations for different rotational angles of the main drive shaft.

[00108] Another possible advantage of the adaptive aperture is the ability to collimate edges that are completely internal to the field. If a treatment plan calls for a volume to be treated that completely surrounds a volume to be protected – for example, a tumor that completely surrounds a spinal cord – a single, machined structure will typically be unable to block radiation to the protected volume without blocking some of treated volume as well. The adaptive aperture can treat such a field using a sequence of leaf positions. For example, the adaptive aperture can be reconfigured dynamically, and during treatment, to protect the areas that require protection, while allowing treatment on areas that require treatment.

[00109] In some cases, better beam performance (penumbra or edge sharpness) results when the particle beam is tangent to the surface of a leaf edge. However, since the beam effectively originates from a single point source, the angle with which it passes through the plane of the adaptive aperture changes as the beam is moved away from the center of the field. For this reason, leaves may have curved edges, as shown in Fig. 1, so that the edges can always be placed a location that makes them tangent to the particle beam. In an example implementation of the adaptive aperture, the tracks on which both primary and secondary carriages move are curved so that flat leaf edges can be used in lieu of curved leaf edges, and so that the flat but remain tangent to the particle beam.

[00110] Fig. 29 shows an example implementation of a curved track 1420. In the example of Fig. 29, the particle beam 1421 originates from source 1422, which may be a particle accelerator like the synchrocyclotron described herein. Particle beam 1422 may scan through field 1423 and, at one extent, may be at location 1425 and, at another extent, may be at location 1426. Carriages holding leaves 1428 and 1429 are mounted on curved track 1420 so that leaves 1428 and 1428 can move towards or away from each other. In this example, the leaves have straight ends (or “fronts”) 1431, 1431 in contrast to curved end 56 of Fig. 1. By employing a curved track, the particle beam can remain tangent, or substantially tangent, to the straight ends throughout the scan field 1423. Keeping the particle beam tangent to the ends

can be advantageous in that it enables the trimming provided by the adaptive aperture to be consistent throughout the range of the beam field.

[00111] To summarize, in some implementations, the adaptive aperture may trim only a small fraction of a treatment area at one time, e.g., an area that is less than the entire treatment area and that is about equal to one spot size, two spot sizes, three spots sizes, four spot sizes, five spot sizes, and so forth. Thus, in some implementations, the adaptive aperture may be small enough to trim a single spot at once and may be large enough to trim several spots in one position, but not the entire field without moving. The adaptive aperture may be configured to move around and within the field as the beam scans. For example, in some implementations, the adaptive aperture tracks the beam as it scans, and its configuration and reconfiguration may be synchronized to the scanning and to the pulses provided by the beam (e.g., different reconfigurations for different beam pulses and/or locations). By not using leaves large enough to trim an entire treatment area, the adaptive aperture can be made smaller and therefore the adaptive aperture can be placed closer to the patient with little or no interference from other devices. In some implementations, no leaf of the adaptive aperture even has one dimension that spans an entire maximum treatment area. In some implementations, each individual leaf is movable in two dimensions within the treatment area, and the device is mounted on a gantry (e.g., in the context of a particle therapy system such as that described herein) to be rotated on one or more axes and made extendable towards and away from isocenter.

[00112] Example treatment planning processes are described in the context of a system, such as that described herein, that contains an adaptive aperture. The example treatment planning processes described below, however, are not limited to use with the systems described herein, and may be used with any appropriate system that applies radiation to a target, such as a tumor, for treatment.

[00113] An example treatment planning process (TPP) performed by a TPS may employ inverse planning. Generally, inverse planning includes obtaining a target dose distribution of radiation, and then performing a process, such as an optimization process, to determine characteristics that together achieve the target

dose distribution. The process may be performed by executing, on one or more processing devices, one or more computer programs that are stored on, and retrieved from, one or more non-transitory machine-readable storage media. The characteristics may include, but are not limited to, spot size, spot shape, spot location, spot weight (the number of protons per spot), spot orientation, and so forth. In some implementations, the characteristics may include information relating to the construction of the adaptive aperture, e.g., the numbers and locations of its leaves, the spot resolution that can be achieved using the adaptive aperture, the ways and extents that the carriages can be moved, and so forth. In such examples, the TPS leverages its knowledge of the structure of the adaptive aperture to produce instructions to the adaptive aperture to control its configuration.

[00114] As described herein, the example adaptive apertures of Figs. 1 to 7, 29, and 30 may be used to trim a single spot, a set of spots, a treatment layer, or a treatment field. A treatment field includes the extent of radiation that can be applied at a particular layer, and may extend beyond the layer itself. An example spot may have a Gaussian distribution, as shown in the graph of Fig. 31. That is, the distribution of protons in the spot is Gaussian in that the center 1500 of the spot has the highest concentration of protons and that concentration decreases as the distance from the center of the spot decreases. This Gaussian distribution can have an effect on treatment planning. For example, in some implementations, a TPP may take this distribution into account in determining the appropriate dose for a given location of a radiation target, and also may take this distribution into account in determining whether, and where, to trim the spot using the adaptive aperture (or any other appropriate spot-trimming mechanism). For example, referring to Fig. 2, the adaptive aperture may be used to trim spots at the border between a radiation target defined by circle 708 (e.g., a tumor) and healthy tissue. However, trimming is not limited to the border between diseased and healthy tissue. For example, a treatment plan may require that spots interior to a treatment layer be trimmed, e.g., in order to achieve a desired dose distribution. Accordingly, the adaptive aperture may be configured to implement interior spot trimming as well.

[00115] An example TPP is executed by a TPS. An example TPP references a library containing information defining characteristics of pre-trimmed spots that the process may use to achieve a target dose distribution. For example, the library may define a trimmed spot based on one or more parameters. For example, referring to Fig. 32, a pre-trimmed spot 1600 (shown in top view) may be defined in terms of its radius and angle of rotation. In the example of Fig. 32, pre-trimmed spot 1600 is created by trimming original spot 1601 at chords 1602 and 1603. In pre-trimmed spot 1600, the radius is the line 1605 from the center 1606 of the non-trimmed spot that is perpendicular to the spot edge. In the example of Fig. 32, the angle of rotation is defined as the amount that the spot has been rotated 1607, either clockwise or counterclockwise, relative to its original position. In some examples, such as spot 1600 of Fig. 32 multiple radii (in this case two) may define the spot. In some implementations, a single radius may define the spot, in part. And, in other implementations, more than two radii may define a spot, in part. The library may store information defining characteristics of any appropriate number of pre-trimmed spots, which may be used to generate actual spots.

[00116] A pre-trimmed spot may have any appropriate shape, and is not limited to the shapes shown in Fig. 32. For example, a pre-trimmed spot may have one or more step-like edges, like the spots trimmed in Figs. 2 and 30.

[00117] In some implementations, additional, or other, information may be used to define spots in the library. For example, area, shape, or other descriptive information may define each spot. This information may be in addition to, or instead of, the radius and rotation indicated above. The information that defines each spot, e.g., a radius and rotation, is stored in computer memory, e.g., in one or more files, look-up tables, or other constructs, that enable the information to be indexed to corresponding leaf movements of the adaptive aperture.

[00118] More specifically, in some implementations, each trimmed spot may be created from a radiation spot (e.g., one generated by the synchrocyclotron described herein) that has been trimmed. As described herein, the trimming may be performed using the adaptive aperture. The configuration of the adaptive aperture used to create each trimmed spot may be indexed to, or otherwise stored in association with,

the information in the library that defines the spot. The configuration may include, but is not limited to, the leaf configuration of the adaptive aperture (e.g., which leaves are extended and which are not), the rotational angle of the adaptive aperture, the position(s) of the primary and secondary carriage(s), and so forth. Any appropriate information that characterizes the configuration of the adaptive aperture to produce the corresponding trimmed spot may be stored.

[00119] In some implementations, during treatment planning, the TPP performs a process, such as an optimization process, in which the TPP selects information for pre-trimmed spots from the library in order to achieve the target dose distribution for the treatment. For example, the TPP may use that information to simulate dose distributions for trimmed spots iteratively until the target dose distribution, or a dose distribution acceptably close to the target dose distribution, is achieved. The information that produces the best dose distribution is received by a computing system that controls the adaptive aperture, and that information is translated into leaf and/or carriage positions. In some implementations, the TPP may employ a Monte Carlo optimization process in order to achieve the desired dose distribution. Generally, a Monte Carlo optimization process includes defining a domain of possible inputs (e.g., information defining trimmed spots from the library), generating inputs randomly from a probability distribution over the domain, performing a deterministic computation on the inputs, and aggregating the results. In some implementations, the TPP may employ a Pencil Beam-based optimization process, or any other appropriate optimization process in order to achieve the desired dose distribution.

[00120] In some implementations, during treatment planning, the TPP performs a process, such as an optimization process, in which the TPP optimizes trimming curves in order to achieve a target dose distribution for the treatment. For example, the TPP may obtain information defining trimming curves to be produced using the adaptive aperture, and use that information to simulate dose distributions for trimmed spots iteratively until the target dose distribution, or a dose distribution acceptably close to the target dose distribution, is achieved. The information defining the trimming curves may include, e.g., lengths, curvatures, angles,

orientations, and any other appropriate information that can define a trimming curve. In this example, the information is independent of the pre-trimmed spots in the library, and the pre-trimmed spots are not used to control the adaptive aperture. Rather, the information that produces the best dose distribution is received by a computing system that controls the adaptive aperture, and that information is translated into leaf and/or carriage positions. As above, in some implementations, the TPP may employ a Monte Carlo optimization process, a Pencil Beam-based optimization process, or any other appropriate optimization process in order to achieve the desired dose distribution.

[00121] In some implementations, during treatment planning, the TPP performs a process, such as an optimization process, in which the TPP defines different spot sizes and/or shapes to be produced through trimming based on information such as, but not limited to, spot radius and orientation. For example, the TPP may use that information to simulate dose distributions for trimmed spots iteratively until the target dose distribution, or a dose distribution acceptably close to the target dose distribution, is achieved. Information identifying the sizes and shapes may include information identifying a spot radius and/or orientation as described herein, or any other appropriate identifying information. In this example, the information is independent of the pre-trimmed spots in the library, and the pre-trimmed spots are not used to control the adaptive aperture. Rather, the information that produces the best dose distribution is received by a computing system that controls the adaptive aperture, and that information is translated into leaf and/or carriage positions. As above, in some implementations, the TPP may employ a Monte Carlo optimization process, a Pencil Beam-based optimization process, or any other appropriate optimization process in order to achieve the desired dose distribution.

[00122] In some implementations, during treatment planning, the TPP performs a process, such as an optimization process, in which the TPP defines leaf and/or carriage movements of the adaptive aperture to perform treatment. That is, the TPP knows the geometry, structures, and capabilities of the adaptive aperture, and leverages this knowledge to optimize the treatment plan. For example, the TPP may use that information to simulate dose distributions for trimmed spots iteratively until

the target dose distribution, or a dose distribution acceptably close to the target dose distribution, is achieved. The leaf and/or carriage positions that are closest to the target dose distribution are provided, by the TPP, to a control/computing system that controls configuration of the adaptive aperture. In this example, the information identifying leaf and/or carriage movements is independent of the pre-trimmed spots in the library, and the pre-trimmed spots are not used to control the adaptive aperture. Rather, the information is received by the control/computing system that controls the adaptive aperture, and that information is used to position the leave(s) and/or carriage(s) of the adaptive aperture. As above, in some implementations, the TPP may employ a Monte Carlo optimization process, a Pencil Beam-based optimization process, or any other appropriate optimization process in order to achieve the desired dose distribution.

[00123] In general, after the TPP determines the appropriate dose distribution, including information to control the adaptive aperture, the TPP provides that information to a control system (e.g., one or more computers) for the particle therapy (or other radiation therapy) system. In some implementations, the information provided from the TPP identifies the spots in the library that are to be used to implement the dose distribution, and indicates where those spots are to be formed on/applied to the irradiation target at the time of delivery of each spot. In some implementations, the information provided from the TPP identifies the spot sizes and shapes that are to be used to implement the dose distribution, and indicates where those spots are to be formed on/applied to the irradiation target at the time of delivery of each spot. Information identifying the sizes and shapes may include information identifying a spot radius and/or orientation as described herein, or any other appropriate identifying information. In some implementations, the information provided from the TPP identifies trimming curves to be produced by the adaptive aperture. In some implementations, the information provided from the TPP includes leaf and/or carriage positions that are specified for the adaptive aperture.

[00124] The control system uses the information received from the TPP to determine the configuration of the adaptive aperture needed to obtain the target dose distribution or a dose distribution that is acceptably close to the target dose

distribution. In an example, information that characterizes the configuration of the adaptive aperture needed to create a trimmed spot may be associated with (e.g., indexed to) information defining the corresponding pre-trimmed spot (e.g., its radius and orientation). The control system obtains the information characterizing the configuration of the adaptive aperture, and uses that information, together with the information indicating where the spots are to be applied at the time of delivery of each spot, to control operation of the particle accelerator and the adaptive aperture. That is, the information indicating where the spots are to be applied at the time of delivery of each spot is used to control beam output and direction; and the information characterizing the configuration of the adaptive aperture is used to configure the adaptive aperture at appropriate times in order to achieve the target dose distribution. In an example, information that characterizes the configuration of the adaptive aperture needed to create a trimming curve may be derived from a shape of that curve received from the TPP. In an example, information that characterizes the configuration of the adaptive aperture may be received directly from the TPP in the form of leaf and/or carriage positions.

[00125] In operation, for each trimmed spot, the adaptive aperture is configured to achieve the appropriate spot size and shape. For example, the leaves of the adaptive aperture are extended or retracted, as appropriate, to produce a trimmed spot that is identical to, or approximates, a spot for a particular location. In some implementations, locations of the adaptive aperture leaves, and therefore movements thereof, are determined relative to a predefined isocenter of the irradiation target. Movements of the leaves are commanded by the control system in accordance with the required configuration for each spot. Movements of the primary and/or secondary carriages may also be commanded, as appropriate, in addition to, or instead of, movements of the leaves. As explained, in some implementations, the size and shape of each spot may vary spot-by-spot (e.g., on a spot-by-spot basis), requiring reconfiguration of the adaptive aperture in between spot scans. In some implementations, the configuration of the adaptive aperture is maintained between spot scans, thereby reducing the need for reconfiguration each time the beam is moved. In some implementations, e.g., at the interior of a layer of an irradiation

target, trimming is not performed, e.g., all of the leaves of the adaptive aperture may be retracted. In some implementations, including at the interior of a layer of the irradiation target, trimming is performed, e.g., to define one or more spot edges to achieve the target dose distribution.

[00126] As described herein, treatment may be performed on the irradiation target layer-by-layer. Layers may be accessed by varying the energy of the particle beam, with higher-energy beams penetrating to reach deeper layers (relative to the beam source). As described herein, the energy of the particle beam may be changed by moving one or more plates into the particle beam path, by changing the energy of the beam using a variable-energy accelerator, or by a combination thereof.

[00127] In some implementations, the adaptive aperture may trim a single spot, part of a layer to be treated, the entire layer to be treated, or the entire radiation field for that layer. That is, the adaptive aperture may be used to trim more than an individual spot or set of spots within a field or layer. For example, the adaptive aperture may be configured differently, for each layer or field for that layer. In this regard, Fig. 33 shows a side view of an irradiation target 1700 having several layers to be treated, including layers 1701 and 1702. Referring also to Figs. 1 to 7, the adaptive aperture may be configured to trim layer 1702; and the adaptive aperture may be configured to trim the layer 1701. In each case, appropriate leaves may be extended or retracted, and one or more carriages may be moved to appropriate locations, in order to change the configuration of the adaptive aperture to a configuration that is appropriate for the layer. The TPP may perform appropriate optimizations to define the trimming to be performed for each layer.

[00128] In some implementations, the adaptive aperture may trim a curve that has been generated by the TPP for all layers of an irradiation target. More specifically, in some implementations, an optimized treatment plan may not require different trimming for different layers (e.g., depths) of an irradiation target. That is, a single trimming curve (an aperture may be a trimming curve) may be configured and used for all layers of the irradiation target, even if those layers have different cross-sectional shapes. For example, in Fig. 34, layers 1801 and 1802 have different widths. The TPP may generate an optimized trimming curve that is capable of

trimming both layers. For example, as shown in Fig. 34, a trimming curve 1804 may trim at edges of wider layers (e.g., 1802) of the irradiation target, and only trim the tail edge 1807 of a Gaussian spot 1806 (depicted as an inverted Gaussian curve) produced at more narrow layers (e.g., 1801). The TPP may generate the optimized trimming curve by dividing the irradiation target into voxels of a predefined shape and volume, and by iteratively adjusting parameters for the voxels until an optimized trimming curve is determined. For example, for each voxel, the TPP selects parameters such as spot weight, location, size, shape, etc., and iterates through those parameters until the target dose distribution is obtained across the irradiation target, or until a dose distribution within a specified tolerance of the target dose distribution is obtained across the irradiation target. After that optimized dose distribution is determined, the TPP may determine a trimming curve that best fits each layer of the irradiation target based on the determined dose distribution.

[00129] More specifically, in an example, the TPP examines the determined dose distribution for each layer and, based on those dose distributions, selects a single trimming curve that is most appropriate for all layers. For example, in some implementations, the trimming curve may be configured to accommodate the irregular patterns of different layers of the irradiation target. Selection of the size and shape of the trimming curve may also be based on structures inside the irradiation target or external to the irradiation target. For example, different exterior structures may be more tolerant of exposure to radiation than others. The trimming curve may be configured to accommodate such tolerances.

[00130] As described, in some implementations, the TPP provides, to the control system, specifications of the trimming curve of Fig. 34. The specifications may include, e.g., the shape and location of the trimming curve relative to a known point, such as an isocenter in the irradiation target. In an example, the control system configures the adaptive aperture to produce a trimming curve that meets the specifications provided by the TPP. In this example, the adaptive aperture is not reconfigured to treat successive layers, and thus mimics a preconfigured (and not adaptive/reconfigurable) aperture or other appropriate structure.

[00131] In the foregoing example, the same trimming curve is used for each layer to be treated. This can be advantageous since it allows the adaptive aperture to be used with legacy TPPs. More specifically, some legacy TPPs are configured for operation with a predefined aperture, where a conventional aperture is a static (e.g., non-adaptive) structure or structures that create a hole through which radiation passes unimpeded. By configuring the adaptive aperture so that it maintains the same configuration throughout treatment, the adaptive aperture can be used with legacy TPPs of the type described above.

[00132] In the example implementations described above, the TPP selects parameters such as spot weight, location, size, shape, etc., and iterates through those parameters until the target dose distribution is obtained across the irradiation target. In some example implementations, the TPP does not have knowledge of, or take into account, the configuration of the adaptive aperture when obtaining the target dose distribution. Rather, the specifications provided by the TPP are interpreted by the control system to control the configuration of the adaptive aperture, e.g., to extend or retract leaves, move carriage(s), etc. in order to achieve an appropriate trimming curve based on the dose distribution.

[00133] As described, in some implementations, the TPP has knowledge of, and takes into account, the configuration of the adaptive aperture when obtaining a target dose distribution. For example, the TPP may know the number of leaves in the adaptive aperture, the dimensions and shape of each leaf, the relative locations of the leaves, the number of carriages, how the carriages move, and so forth. This information may be taken into account, along with, e.g., the spot weight, spot location, spot size, spot orientation, and/or spot shape, when determining an optimized dose distribution for the irradiation target. The TPP may use this information in determining appropriate leaf positions corresponding to each delivered spot. That is, the TPP has knowledge of the true geometry and motion constraints of the adaptive aperture leaves and includes that geometry directly in its optimization calculations. The TPP then produces a prescription for the exact positions of the leaves and/or carriages for every spot. Thus, in this example, the specifications provided by the TPP to the control system include information used to configure the

adaptive aperture. For example, the specifications may include a configuration of leaves and/or carriages of the adaptive aperture. In some implementations, a different set of configurations of the adaptive aperture, as appropriate, may be provided for each layer of, or spot on, the irradiation target. In this example, information about the trimming curve itself need not be provided by the TPP, since the configuration information for the adaptive aperture (including, e.g., the leaf configuration information) provided by the TPP defines the trimming to be produced.

[00134] In an example implementation, the TPP performs an optimization process, such as the Monte Carlo optimization or the Pencil Beam optimization, using information defining different spot shapes (e.g., obtained from the library or elsewhere) and using known adaptive aperture (e.g., leaf) configurations. In an example optimization process, for a known particle beam, the TPP may iterate through different combinations of beam shape and adaptive aperture configurations to obtain a resulting dosage, and compare the resulting dosage to a target dosage (for the area under consideration). If the resulting dosage deviates from the target dosage by more than a defined amount, a different combination of spot shape and adaptive aperture configuration is selected, and the process is repeated until an optimization is achieved (e.g., the resulting dosage matches the target dosage for the area or is within the defined tolerance relative to the target dosage).

[00135] In some implementations, a TPP selects beam spot shapes based on the dose and dose restrictions of an inverse planning process. In some implementations, the selection of spot shapes and sizes may be based on spots defined in the library described herein that contains information defining characteristics of pre-trimmed spots. In some implementations, the selection of spot shapes and sizes may be independent of spots defined in the library described herein that contains information defining characteristics of pre-trimmed spots. For example, TPP may configure the adaptive aperture – including leaves and/or carriages – to produce spots having any appropriate shapes. In other words, the TPP may be programmed with the capabilities of the particle therapy system, and also be programmed with a treatment plan for an irradiation target, such as a tumor. Given this information, the TPP may instruct the adaptive aperture to produce spot

shapes and sizes by trimming the particle beam in any appropriate manner to achieve the desired treatment. The trimming that is produced by the adaptive aperture is not limited to producing predefined spots from the library in this example.

[00136] In some implementations, a TPP selects beam spot shapes based on the dose and dose restrictions of an inverse planning process. A radiation plan that describes the positioning and shape of multiple treatment spots may be developed with reference to a dose map prepared by a physician using planning software to convert the dose map to a radiation plan. The dose map may be prepared, for example, using a graphics terminal with the physician viewing one or more CT images of the patient to define target (e.g., desired) doses in different zones within the volume of the patient. An optimization process may be performed to determine which spots to apply to the treatment area to achieve the target dose.

[00137] Fig. 37 shows an example of an optimization process 2100 that may be performed, e.g., by the TPP. The process may be implemented as a computer program stored on one or more non-transitory machine-readable media. The instructions may be executed on one or more processing devices (e.g., microprocessors) to implement the operations shown in Fig. 37. In this regard, although the process of Fig. 37 is defined in terms of physical entities, such as spots, intensities, treatment areas, and so forth, the process of Fig. 37 is not performed on those physical entities, but rather is performed using data that represents those physical entities and that models those physical entities and their interactions. Accordingly, although the description of Fig. 37 and the accompanying figures refers to the entities themselves (e.g., spots, treatment areas, etc.), the entity names are actually shorthand for data describing those entities (as the modeling and optimization process are performed on a computer, and output a treatment plan that is later implemented using a therapy system, such as those described herein).

[00138] The optimization process includes optimizing (2101) a coarse beam set. Referring to Fig. 38, in this example, this operation (2101) includes obtaining a minimum objective dosage 2202 and a maximum objective dosage 2201. In an example, the minimum objective is the minimum dosage to be applied to a treatment area 2204, such as a tumor (e.g., 98% of a desired dose over 98% of the treatment

area). The maximum objective is the maximum amount of dosage to be applied outside the treatment area, e.g., to normal tissue or critical structures, such as organs. In some implementations, the dosage may be for a single scan (e.g., part of a target cumulative dose) and in some implementations the dosage may be the target cumulative dose.

[00139] The minimum dosage and maximum dosage may be determined, e.g., by a physician in the manner described above based on various factors. Data defining a set of Gaussian spots (in this example, spots A, B, C, D, and E of Fig. 39) is used to model the effect of radiation in those spots on the treatment area. Spots A, B, C, D, and E constitute an initial estimation of which spots will provide treatment to the treatment area given the constraints of the minimum objective and the maximum objective, e.g., spots that come closest to achieving the minimum objective but not exceeding the maximum objective. However, at this point, those objectives may not actually be met. At this stage – the coarse beam set – the spots have the same, or about the same, intensity, as depicted graphically in Fig. 39. In some implementations, different sets of coarse spots may be selected and the results of their application modeled. The set that most closely approximates the target dosage may be the set that is selected as the coarse set. For example, a set that is within a defined percentage of the target dose distribution may be selected, with the target percentage being any appropriate percentage, e.g., 5%, 10%, 15%, 20%, etc.

[00140] The optimization process includes performing (2102) spot intensity optimization to further approximate the target dosage. In this example, spot intensity optimization includes identifying which spots should be increased or decreased in intensity given constraints, such as the minimum objective and the maximum objective (based, e.g., on the location of the treatment area, the location of normal tissue, the location of critical structures, and other appropriate criteria). Referring to Figs. 39 and 40, in this example, the optimization process recognizes that spot E is at the edge of the treatment area 2204 and mostly outside of the treatment area. Therefore, the optimization process decreases the intensity of spot E (e.g., to some level at or below the maximum objective). This is depicted graphically in Fig. 40 by

spot E being reduced in size and/or shape relative to spots A, B, C, and D. This optimization may be performed for multiple spots, as appropriate, and brings the modeled dosage closer to the target dosage. In some implementations, different beam intensities may be selected and the results of their application modeled. The beam intensity that most closely approximates the target dosage at the appropriate location may be the beam intensity that is selected.

[00141] The optimization process includes identifying (2103) spots within (e.g., that affect) one or more transition regions in the treatment area. For example a transition region may be the edge 2205 of the treatment area, and a spot at the transition region may include spot D (as evidenced by line 2206). Notably, a transition region need not only be at an edge of the treatment area. As described herein, transition regions may be at locations internal to the treatment area, and may depend upon the dose profile, the locations of critical structures in the treatment area, the shape of a tumor to be treated, and other appropriate factors.

[00142] The optimization process includes replacing (2104) a spot (or spots) at the transition region(s) with one or more spots of alternate shapes. In some implementations, referring to Fig. 41, this operation includes replacing spot D at the transition regions with a pre-trimmed spot from the library that best approximates the target dose distribution. For example, Fig. 42 shows side views of examples of pre-trimmed spots in the library (with the side views showing distributions of protons within the spot, e.g., Gaussian or trimmed from Gaussian). An appropriate spot, such as spot 2401, is selected and replaced for spot D, with resulting replaced spot D' having (in this case) the shape shown in Fig. 41. In some implementations, differently-shaped spots may be selected from the library and the results of their application modeled. The pre-trimmed spot that most closely approximates the target dosage at the appropriate location may be selected for use. Operations 2101 through 2104 may be repeated, as necessary and appropriate, for different treatment areas and different locations within a single treatment area.

[00143] In some implementations, operation 2104 does not include use of pre-trimmed spots from the library. Instead, operation 2104 may include defining spot shapes based on information including, but not limited to, the desired treatment that

is to be applied to an irradiation target and the ways in which the leaves and carriages of the adaptive aperture can move to produce one or more trimming curves. The configuration of one or more of the spots may be defined in real-time and, in some implementations, are not predefined. The spots may be defined using radius and/or orientation information for each spot, alone or in combination with any other appropriate information that can be used to specify spot size and shape. In this regard, in some implementations, real-time may include actions that occur on a continuous basis or track each other in time, taking into account delays associated with processing, data transmission, hardware, and the like.

[00144] The optimization process outputs (2105) information that defines which spots to apply to the treatment area. Some of the spots may not be trimmed (e.g., spots A, B, C), in which case only intensity and other appropriate information is provided to the therapy control system. For trimmed spots (e.g., spot D), those match the spots from the library, and information identifying those pre-trimmed spots is output to the control system. For implementations that do not employ pre-trimmed spots, information characterizing spot size and shape may be used. As described herein, information in the library defining the shape of those spots is indexed or otherwise associated with configurations of the adaptive aperture. Accordingly, that information is obtained by the control system, and used to control the adaptive aperture, as appropriate, to form those spots. For intensity-modulated spots, information for those spots is also output and used to produce spots with appropriate intensity. In some implementations, intensity-modulates spots may also be pre-trimmed spots that have been modulated in accordance with a desired intensity.

[00145] As noted above, a trimming curve may define an aperture, e.g., one or more curves or a hole that limit the extent to which radiation passes to an irradiation target. In some implementations, a trimming curve may simply be a line or curve that defines an edge of a spot or field of an irradiation target. For example referring to Fig. 35, a layer 1901 may include multiple internal trimming curves 1902 to 1905, none of which are closed curves.

[00146] Fig. 36 shows how configuration of the adaptive aperture fits into the process of Fig. 17. Other than operations 2001 and 2002, the operations of Fig. 36

are the same as those of Fig. 17. In this example, as part of the layer selection process or prior thereto, the control system of a particle therapy system receives (2001), from a TPP, information that is based on a dose distribution for an irradiation target; and configures (2002) the adaptive aperture by performing at least one of the following operations: (A) moving structures to trim spots of a particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received; (B) moving structures to produce a trimming curve for each layer of an irradiation target based on specifications of trimming curves for different layers included in the information received; (C) moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included in the information received; or (D) moving structures to produce a trimming curve for a radiation field of an irradiation target based on configuration information for the structures in the information received.

[00147] As described herein, in some implementations, the structures may be leaves that are moved into, or out of, the beam path. In some implementations, process 2000 of Fig. 36 may be implemented with a device other than the adaptive aperture. As appropriate, any two or more of the foregoing operations (A), (B), (C), or (D) may be combined, or used together, in whole or in part.

[00148] Also described herein is an example of a particle accelerator system, such as a proton or ion therapy system, which may employ adaptive apertures of Figs. 1 to 7, 29, and 30, and which may configure those adaptive apertures in accordance with information from a TPP as described herein. The example particle therapy system includes a particle accelerator – in this example, a synchrocyclotron – mounted on a gantry. The gantry enables the accelerator to be rotated around a patient position, as explained in more detail below, to allow a particle beam from the particle accelerator hit any arbitrary treatment area in the patient. In some implementations, the gantry is steel and has two legs mounted for rotation on two respective bearings that lie on opposite sides of a patient. The particle accelerator is supported by a steel truss that is long enough to span a treatment area in which the patient lies and that is attached at both ends to the rotating legs of the gantry. As a

result of rotation of the gantry around the patient, the particle accelerator also rotates.

[00149] In an example implementation, the particle accelerator (e.g., the synchrocyclotron) includes a cryostat that holds one or more superconducting coils, each for conducting a current that generates a magnetic field (B). In an example, the cryostat uses liquid helium (He) to maintain each coil at superconducting temperatures, e.g., 4° Kelvin (K). Magnetic yokes or smaller magnetic pole pieces are located inside the cryostat, and define the shape of a cavity in which particles are accelerated. Magnetic shims may pass through the magnetic yokes or pole pieces to change the shape and/or magnitude of the magnetic field in the cavity.

[00150] In this example implementation, the particle accelerator includes a particle source (e.g., a Penning Ion Gauge – PIG source) to provide an ionized plasma column to the cavity. Hydrogen gas is ionized to produce the plasma column. A voltage source provides a radio frequency (RF) voltage to the cavity to accelerate pulses of particles from the plasma column into the cavity. The magnetic field in the cavity is shaped to cause particles to move orbitally within the cavity. The magnetic field may be, e.g., at least 4 Tesla, as explained herein.

[00151] As noted, in an example, the particle accelerator is a synchrocyclotron. Accordingly, the RF voltage is swept across a range of frequencies to account for relativistic effects on the particles (e.g., increasing particle mass) when accelerating particles from the plasma column. The magnetic field produced by running current through a superconducting coil, together with the shape of the cavity, causes particles accelerated from the plasma column to accelerate orbitally within the cavity. In other implementations, a particle accelerator other than a synchrocyclotron may be used. For example, a cyclotron, a synchrotron, a linear accelerator, and so forth may be substituted for the synchrocyclotron described herein.

[00152] In the example synchrocyclotron, a magnetic field regenerator (“regenerator”) is positioned near the outside of the cavity (e.g., at an interior edge thereof) to adjust the existing magnetic field inside the cavity to thereby change locations (e.g., the pitch and angle) of successive orbits of the particles accelerated from the plasma column so that, eventually, the particles output to an extraction

channel that passes through the cryostat. The regenerator may increase the magnetic field at a point in the cavity (e.g., it may produce a magnetic field “bump” of about 2 Tesla or so at an area of the cavity), thereby causing each successive orbit of particles at that point to precess outwardly toward the entry point of the extraction channel until it reaches the extraction channel. The extraction channel receives, from the cavity, particles accelerated from the plasma column and outputs the received particles from the cavity as a particle beam.

[00153] The superconducting (“main”) coils can produce relatively high magnetic fields. The maximum magnetic field generated by a main coil may be within a range of 4T to 20T or more. For example, a main coil may be used to generate magnetic fields at, or that exceed, one or more of the following magnitudes: 4.0T, 4.1T, 4.2T, 4.3T, 4.4T, 4.5T, 4.6T, 4.7T, 4.8T, 4.9T, 5.0T, 5.1T, 5.2T, 5.3T, 5.4T, 5.5T, 5.6T, 5.7T, 5.8T, 5.9T, 6.0T, 6.1T, 6.2T, 6.3T, 6.4T, 6.5T, 6.6T, 6.7T, 6.8T, 6.9T, 7.0T, 7.1T, 7.2T, 7.3T, 7.4T, 7.5T, 7.6T, 7.7T, 7.8T, 7.9T, 8.0T, 8.1T, 8.2T, 8.3T, 8.4T, 8.5T, 8.6T, 8.7T, 8.8T, 8.9T, 9.0T, 9.1T, 9.2T, 9.3T, 9.4T, 9.5T, 9.6T, 9.7T, 9.8T, 9.9T, 10.0T, 10.1T, 10.2T, 10.3T, 10.4T, 10.5T, 10.6T, 10.7T, 10.8T, 10.9T, 11.0T, 11.1T, 11.2T, 11.3T, 11.4T, 11.5T, 11.6T, 11.7T, 11.8T, 11.9T, 12.0T, 12.1T, 12.2T, 12.3T, 12.4T, 12.5T, 12.6T, 12.7T, 12.8T, 12.9T, 13.0T, 13.1T, 13.2T, 13.3T, 13.4T, 13.5T, 13.6T, 13.7T, 13.8T, 13.9T, 14.0T, 14.1T, 14.2T, 14.3T, 14.4T, 14.5T, 14.6T, 14.7T, 14.8T, 14.9T, 15.0T, 15.1T, 15.2T, 15.3T, 15.4T, 15.5T, 15.6T, 15.7T, 15.8T, 15.9T, 16.0T, 16.1T, 16.2T, 16.3T, 16.4T, 16.5T, 16.6T, 16.7T, 16.8T, 16.9T, 17.0T, 17.1T, 17.2T, 17.3T, 17.4T, 17.5T, 17.6T, 17.7T, 17.8T, 17.9T, 18.0T, 18.1T, 18.2T, 18.3T, 18.4T, 18.5T, 18.6T, 18.7T, 18.8T, 18.9T, 19.0T, 19.1T, 19.2T, 19.3T, 19.4T, 19.5T, 19.6T, 19.7T, 19.8T, 19.9T, 20.0T, 20.1T, 20.2T, 20.3T, 20.4T, 20.5T, 20.6T, 20.7T, 20.8T, 20.9T, or more. Furthermore, a main coil may be used to generate magnetic fields that are within the range of 4T to 20T (or more, or less) that are not specifically listed above.

[00154] In some implementations, such as the implementation shown in Figs. 8 and 9, large ferromagnetic magnetic yokes act as a return for stray magnetic field produced by the superconducting coils. For example, in some implementations, the superconducting magnet can generate a relatively high magnetic field of, e.g., 4T or

more, resulting in considerable stray magnetic fields. In some systems, such as that shown in Figs. 8 and 9, the relatively large ferromagnetic return yoke 100 is used as a return for the magnetic field generated by superconducting coils. A magnetic shield surrounds the yoke. The return yoke and the shield together dissipated stray magnetic field, thereby reducing the possibility that stray magnetic fields will adversely affect the operation of the accelerator.

[00155] In some implementations, the return yoke and shield may be replaced by, or augmented by, an active return system. An example active return system includes one or more active return coils that conduct current in a direction opposite to current through the main superconducting coils. In some example implementations, there is an active return coil for each superconducting coil, e.g., two active return coils – one for each superconducting coil (referred to as a “main” coil). Each active return coil may also be a superconducting coil that surrounds the outside of a corresponding main superconducting coil concentrically.

[00156] Current passes through the active return coils in a direction that is opposite to the direction of current passing through the main coils. The current passing through the active return coils thus generates a magnetic field that is opposite in polarity to the magnetic field generated by the main coils. As a result, the magnetic field generated by an active return coil is able to dissipate at least some of the relatively strong stray magnetic field resulting from the corresponding main coil. In some implementations, each active return may be used to generate a magnetic field of between 2.5T and 12T or more. For example, the magnetic field may be 2.5T, 2.6T, 2.7T, 2.8T, 2.9T, 3.0T, 3.1T, 3.2T, 3.3T, 3.4T, 3.5T, 3.6T, 3.7T, 3.8T, 3.9T, 4.0T, 4.1T, 4.2T, 4.3T, 4.4T, 4.5T, 4.6T, 4.7T, 4.8T, 4.9T, 5.0T, 5.1T, 5.2T, 5.3T, 5.4T, 5.5T, 5.6T, 5.7T, 5.8T, 5.9T, 6.0T, 6.1T, 6.2T, 6.3T, 6.4T, 6.5T, 6.6T, 6.7T, 6.8T, 6.9T, 7.0T, 7.1T, 7.2T, 7.3T, 7.4T, 7.5T, 7.6T, 7.7T, 7.8T, 7.9T, 8.0T, 8.1T, 8.2T, 8.3T, 8.4T, 8.5T, 8.6T, 8.7T, 8.8T, 8.9T, 9.0T, 9.1T, 9.2T, 9.3T, 9.4T, 9.5T, 9.6T, 9.7T, 9.8T, 9.9T, 10.0T, 10.1T, 10.2T, 10.3T, 10.4T, 10.5T, 10.6T, 10.7T, 10.8T, 10.9T, 11.0T, 11.1T, 11.2T, 11.3T, 11.4T, 11.5T, 11.6T, 11.7T, 11.8T, 11.9T, 12.0T, or more. Furthermore, an active return coil may be used to generate

magnetic fields that are within the range of 2.5T to 12T (or more, or less) that are not specifically listed above.

[00157] Referring to Fig. 10, at the output of extraction channel 102 of particle accelerator 105 (which may have the configuration shown in Figs. 8 and 9), is an example scanning system 106 that may be used to scan the particle beam across at least part of an irradiation target. Fig. 11 also shows examples of components of the scanning system. These include, but are not limited to, a scanning magnet 108, an ion chamber 109, and an energy degrader 110. Other components that may be incorporated into the scanning system are not shown in Fig. 11, including, e.g., one or more scatterers for changing beam spot size. The example scanning system, including components thereof, may be mounted to the gantry and move along with the particle accelerator during movement of the gantry.

[00158] In an example operation, scanning magnet 108 is controllable in two dimensions (e.g., Cartesian XY dimensions) to direct the particle beam across a treatment area (e.g., a cross-section) of an irradiation target. Ion chamber 109 detects the dosage of the beam and feeds-back that information to a control system to adjust beam movement. Energy degrader 110 is controllable to move material (e.g., one or more individual plates) into, and out of, the path of the particle beam to change the energy of the particle beam and therefore the depth to which the particle beam will penetrate the irradiation target. In this way, the energy degrader selects a depth-wise layer of an irradiation target to scan in two dimensions.

[00159] Figs. 12 and 13 show views of an example scanning magnet 108. In this example implementation, scanning magnet 108 includes two coils 111, which control particle beam movement in the X-dimension, and two coils 112, which control particle beam movement in the Y-dimension. Control is achieved, in some implementations, by varying current through one or both sets of coils to thereby vary the magnetic field(s) produced thereby. By varying the magnetic field(s) appropriately, the particle beam can be moved in the X- and/or Y- dimension across the irradiation target. In some implementations, the scanning magnet is not movable physically relative to the particle accelerator. In other implementations, the scanning magnet may be movable relative to the particle accelerator (e.g., in addition to the

movement provided by the gantry). In some implementations, the scanning magnet may be controllable to move the particle beam continuously so that there is uninterrupted motion of the particle beam over at least part of, and possibly all of, a layer of an irradiation target being scanned or a portion thereof (e.g., a treatment area). In other implementations, the scanning magnets are controllable at intervals or specific times. In some implementations, there may be different scanning magnets to control all or part movement of a particle beam in the X- and/or Y-dimensions.

[00160] In some implementations, scanning magnet 108 may have an air core. In other implementations, scanning magnet 108 may have a ferromagnetic (e.g., an iron) core. In general, a magnet having an air core includes a magnetic coil around a core that is a non-ferromagnetic material, such as air. For example, an air core magnet may include self-supporting coils that surround air. In some implementations, an air core magnet may include coils that are wound around an insulator, such as ceramic or plastic, which may or may not include air.

[00161] In some cases, an air core may have advantages over a ferromagnetic core. For example, the amount that the particle beam moves (e.g., is deflected) in the X- and/or Y-dimensions is determined, at least in part, based on the amount of current applied to the magnet (referred to as the "magnet current"). A scanning magnet typically has a movement (or deflection) range, which is the extent over which the magnet will move the beam. At extremes of this range, such as at the edges, larger amounts of current are applied to the scanning magnet in order to achieve relatively high amounts of beam deflection. Some types of scanning magnets having a ferromagnetic core may saturate at these extremes, resulting in a non-linear relationship between current and magnet movement. That is, the amount of deflection produced by the magnet may not be linearly proportional to the amount of current applied to the magnet. Due to this non-linearity, in some cases, it may be difficult to determine and/or set some beam locations using magnet current. Accordingly, when a scanning magnet having a ferromagnetic core is used, there may need to be some calibration and/or compensation performed in order to correct for non-linearities such as that described above.

[00162] In contrast, a scanning magnet having an air core may not saturate in the same manner as a scanning magnet having a ferromagnetic core. For example, an air core magnet may not saturate or may saturate less than a magnet having a ferromagnetic core. As a result, the relationship between current and magnet movement may be more linear, particularly at the range extremes, making determinations of beam location based on magnet current more accurate, at least in some cases. This increased linearity also can enable more accurate movement of the beam, particularly at range extremes. That is, since the relationship between current and beam movement is generally more linear over a larger range when an air core scanning magnet is used, beam movement may be more easily reproducible using an air core scanning magnet. This can be advantageous, since a depth-wise layer of an irradiation target may require multiple scans, each providing a percentage of a total cumulative radiation dose. Precision in delivery of multiple doses to the same area, such as that which can be obtained through use of an air core scanning magnet, can affect the efficacy of the treatment.

[00163] Although the relationship between current and magnet movement may be more linear in an air core magnet, in some cases, an air core magnet may be more susceptible to stray magnetic fields than a magnet having a ferromagnetic core. These stray magnetic fields may impact the scanning magnet during motion of the scanning magnet produced by the gantry. Accordingly, in some implementations that use a scanning magnet having an air core, the current applied to the scanning magnet to move the beam may be calibrated to account for the position of the scanning magnet relative to the patient (or, correspondingly, to account for the position of the gantry, since the position of the gantry corresponds to the position of the scanning magnet relative to the patient). For example, the behavior of the scanning magnet may be determined and, if necessary, corrected, for different rotational positions (angles) of the gantry, e.g., by increasing or decreasing some applied current based on rotational position.

[00164] In some implementations, the scanning magnet may have a core that is comprised of both air and a ferromagnetic material (e.g., iron). In such

implementations, the amount and configuration of air and ferromagnetic material in the core may be determined taking the foregoing factors into account.

[00165] In some implementations, a current sensor 118 may be connected to, or be otherwise associated with, scanning magnet 108. For example, the current sensor may be in communication with, but not connected to, the scanning magnet. In some implementations, the current sensor samples current applied to the magnet, which may include current to coil(s) for controlling beam scanning in the X-dimension and/or current to coil(s) for controlling beam scanning in the Y-dimension. The current sensor may sample current through the magnet at times that correspond to the occurrence of pulses in the particle beam or at a rate that exceeds the rate that the pulses occur in the particle beam. In the latter case, the samples, which identify the magnet current, are correlated to detection of the pulses by the ion chamber described below. For example, the times at which pulses are detected using the ion chamber (described below) may be correlated in time to samples from the current sensor, thereby identifying the current in the magnet coil(s) at the times of the pulses. Using the magnet current, it thus may be possible to determine the location on the irradiation target (e.g., on a depth-wise layer of the irradiation target) where each pulse, and thus dose of particles, was delivered. The location of the depth-wise layer may be determined based on the position of the energy degrader (e.g., the number of plates) in the beam path.

[00166] During operation, the magnitude(s) (e.g., value(s)) of the magnet current(s) may be stored for each location at which a dose is delivered, along with the amount (e.g., intensity) of the dose. A computer system, which may be either on the accelerator or remote from the accelerator and which may include memory and one or more processing devices, may correlate the magnet current to coordinates within the irradiation target, and those coordinates may be stored along with the amount of the dose. For example, the location may be identified by depth-wise layer number and Cartesian XY coordinates or by Cartesian XYZ coordinates (with the layer corresponding to the Z coordinate). In some implementations, both the magnitude of the magnet current and the coordinate locations may be stored along with the dose at each location. This information may be stored in memory either on,

or remote from, the accelerator. As described in more detail herein, this information may be used during scanning to apply multiple doses to the same locations to achieve target cumulative doses.

[00167] In some implementations, ion chamber 109 detects dosage (e.g., one or more individual doses) applied by the particle beam to positions on an irradiation target by detecting the numbers of ion pairs created within a gas caused by incident radiation. The numbers of ion pairs correspond to the dose provided by the particle beam. That information is fed-back to the computer system and stored in memory along with the time that the dose is provided. This information may be correlated to, and stored in association with, the location at which the dose was provided and/or the magnitude of the magnet current at that time, as described above.

[00168] As described in more detail below, in some implementations, the scanning system is run open loop, in which case the particle beam is moved freely and uninterrupted across an irradiation target so as to substantially cover the target with radiation. As the radiation is delivered, dosimetry implemented by the particle therapy control system records (e.g., stores) the amount of the radiation per location and information corresponding to the location at which the radiation was delivered. The location at which the radiation was delivered may be recorded as coordinates or as one or more magnet current values, and the amount of the radiation that was delivered may be recorded as dosage in grays. Because the system is run open loop, the delivery of the radiation is not synchronized to the operation of the particle accelerator (e.g., to its RF cycle). However, the dosimetry may be synchronized to the operation of the particle accelerator. More specifically, the dosimetry records the amount and location of each dose delivered as the dose is delivered (that is, as close in time to delivery as possible given the limits of technology). Since the dose is delivered in synchronism with the operation of the accelerator (e.g., one pulse is delivered per RF cycle), in some implementations, the dosimetry that records the dose and the location operates in synchronism, or substantially in synchronism, with delivery of radiation doses to the target, and thus in synchronism with the operation of the particle accelerator, such as its RF cycle.

[00169] Fig. 14 shows a range modulator 115, which is an example implementation of energy degrader 110. In some implementations, such as that shown in Fig. 14, the range modulator includes a series of plates 116. The plates may be made of one or more of the following example materials: carbon, beryllium or other material of low atomic number. Other materials, however, may be used in place of, or in addition to, these example materials.

[00170] One or more of the plates is movable into, or out of, the beam path to thereby affect the energy of the particle beam and, thus, the depth of penetration of the particle beam within the irradiation target. For example, the more plates that are moved into the path of the particle beam, the more energy that will be absorbed by the plates, and the less energy the particle beam will have. Conversely, the fewer plates that are moved into the path of the particle beam, the less energy that will be absorbed by the plates, and the more energy the particle beam will have. Higher energy particle beams typically penetrate deeper into the irradiation target than do lower energy particle beams. In this context, “higher” and “lower” are meant as relative terms, and do not have any specific numeric connotations.

[00171] Plates are moved physically into, and out of, the path of the particle beam. For example, as shown in Fig. 15, a plate 116a moves along the direction of arrow 117 between positions in the path of the particle beam and outside the path of the particle beam. The plates are computer-controlled. Generally, the number of plates that are moved into the path of the particle beam corresponds to the depth at which scanning of an irradiation target is to take place. For example, the irradiation target can be divided into cross-sections or depth-wise layers, each of which corresponds to an irradiation depth. One or more plates of the range modulator can be moved into, or out of, the beam path to the irradiation target in order to achieve the appropriate energy to irradiate each of these cross-sections or depth-wise layers of the irradiation target. The range modulator may be stationary relative to the particle beam during scanning of a part of (e.g., cross-section of) an irradiation target, except for its plates moving in and out of the path of the particle beam. Alternatively, the range modulator of Figs. 14 and 15 may be replaced with a range

modulator that, at least some of the time, tracks movement of the particle beam, thereby enabling use of smaller plates.

[00172] In implementations that use a range modulator of the type described above, the number of plates that are moved into the beam path determine/set the depth-wise layer of the irradiation target that is to be scanned. For example, if two plates are moved into the beam path, the layer will be more shallow than if one or no plates are moved into the beam path. The layer may be identified, and stored in memory, based on the number of plates moved into the beam path. In some implementations, the plates may have different heights. In such implementations, the heights of the various plates also affect which layer is to be scanned (e.g., how deep the particle beam will penetrate the target).

[00173] In some implementations, the particle accelerator may be a variable-energy particle accelerator. In example systems where a variable-energy particle accelerator is used, there may be less need for an energy degrader of the type described herein, as the energy level of the particle beam may be controlled by the particle accelerator. For example, in some systems that employ a variable-energy particle accelerator, an energy degrader may not be needed. In some systems that employ a variable-energy particle accelerator, an energy degrader may still be used to change beam energy levels.

[00174] In some implementations, a treatment plan is established prior to treating the irradiation target. The treatment plan may be stored in memory that is accessible to a computer system that controls operation of the particle therapy system. As described, the treatment plan may be received from a TPS/TPP and may include information about how radiation treatment is to be provided by the particle therapy system. For example, the treatment plan may specify how scanning is to be performed for a particular irradiation target. In some implementations, the treatment plan specifies that raster scanning is to be performed. Raster scanning includes producing an uninterrupted movement of the particle beam across the irradiation target. For example, the scanning magnet moves continually to scan (e.g., move) the particle beam across the irradiation target so as to produce uninterrupted movement of the particle beam over at least part of a layer of an

irradiation target. The movement may be uninterrupted across an entire layer of the irradiation target or across only part of a layer. In some implementations, the beam may be moved at a constant speed along all or part of a layer of the irradiation target. In some implementations, the speed at which the beam is moved along all or part of a layer of the irradiation target may vary. For example, the particle beam may move more quickly across internal portions of a layer than at edges of the layer. The speed of movement may be specified in the treatment plan.

[00175] In some implementations, the treatment plan may also specify the target cumulative dose of radiation (particles) to be applied to various positions on layers of an irradiation target. The dose is cumulative in the sense that it may be achieved through application of one or more doses of particles. For example, the same location (e.g., in XYZ space) on an irradiation target may be irradiated ten times, each time with 10% of the target cumulative dose to achieve the target cumulative dose. In some implementations, the treatment plan need not specify the amount of dose for each location, the locations, or the number of times that locations are to be irradiated. That is, this information may be omitted from the treatment plan in some implementations. Rather, in some implementations, the intensity of the particle beam may be set beforehand to provide a particular dose of radiation per instance of irradiation. The particle beam may then be scanned over a layer of the irradiation target in an open loop manner, without requiring feedback to move to a next location. As the particle beam is scanned, the location of the beam is determined and the corresponding dose at that location is determined. The determination may be made at about the same time as the scanning and delivery (that is, as close in time to delivery as possible given the limits of technology). The cumulative dose at that location, which includes the current dose as well as any dose previously delivered during the current treatment, is compared to the target cumulative dose from the treatment plan. If the two do not match, then additional dose may be applied to that location during a subsequent scan. Since it is not always known precisely how much radiation will be delivered to a location per scan, the number of times that a location is scanned may not be set beforehand. Likewise, since there may be fluctuations in the amount of radiation actually

delivered per scan to a location, the precise amount of radiation per scan is not necessarily set beforehand. Consequently, in some implementations, such information need not be included in the treatment plan.

[00176] In some implementations, the treatment plan may also include one or more patterns, over which the particle beam may be scanned per layer. The treatment plan may also specify the number of plates of an energy degrader to achieve a particular energy level/layer. Other implementations may include information in addition to, or instead of, that specified above.

[00177] In some implementations, the overall treatment plan of an irradiation target may include different treatment plans for different cross-sections (layers) of the irradiation target. The treatment plans for different cross-sections may contain the same information or different information, such as that provided above.

[00178] The scanning system may include the adaptive aperture of Figs. 1 to 7, 29, and 30 (or a variant thereof), which is placeable relative to the irradiation target to limit the extent of the particle beam and thereby limit the extent of the particle beam. For example, the adaptive aperture may be placed in the beam path down-beam of the energy degrader and before the particle beam hits the treatment area of the irradiation target. The adaptive aperture is controllable to allow at least part of the particle beam to pass therethrough and then hit certain parts of the treatment area, while preventing the particle beam from hitting other parts of the patient. For example, the adaptive aperture may be controlled to prevent the particle beam from hitting healthy tissue or to prevent the particle beam from hitting other parts of the irradiation target (e.g., if certain parts of the target are to receive more radiation than other parts). The adaptive aperture may be controlled to perform trimming internal to, or external to, a layer of an irradiation target, as described. Fig. 16 depicts placement of an implementation 770 of the adaptive aperture described herein relative to a patient 771. The direction of beam 771a is also shown.

[00179] As noted above, in some implementations, scanning is performed in an open-loop manner, e.g., by an open-loop control system that may be implemented using one or more processing devices, such as the computing device that controls the particle therapy system. In this example, open-loop scanning includes moving

the particle beam across an irradiation target to substantially cover the target with radiation. In some implementations, movement is not synchronized with operation of the accelerator, e.g., with the RF frequency, but rather runs independently of the operation of the accelerator when the accelerator is operating. For example, movement of the particle beam may be uninterrupted, and not dependent upon the RF cycle of the particle accelerator. Uninterrupted movement may occur across all or part of a layer of an irradiation target. However, as described herein, the dosimetry may be synchronized with delivery of pulses of the particle beam to the irradiation target. In examples where the dosimetry is synchronized with delivery of pulses of the particle beam, the dosimetry is also synchronized with operation of the accelerator (e.g., with the RF frequency used to draw pulses of the particle beam from the ion source plasma column).

[00180] The radiation level of an individual dose of particle beam (e.g., an individual pulse from the accelerator) may be set beforehand. For example, each individual dose may be specified in grays. An individual dose may be, or correspond to, a percentage of the target cumulative dose that is to be applied to a location (e.g., an XYZ coordinate) in an irradiation target. In some implementations, the individual dose may be 100% of the target cumulative dose and, as a result, only a single scan may be needed to deliver a single dose of radiation (e.g., one or more particle pulses) per location to the irradiation target. In some implementations, the individual dose may be less than 100% of the target cumulative dose, resulting in the need for multiple scans of the same location to deliver multiple doses of radiation to the irradiation target. The individual dose may be any appropriate percentage of the target cumulative dose, such as: 1%, 2%, 3%, 4%, 5%, 6%, 7%, 8%, 9%, 10%, 11%, 12%, 13%, 14%, 15%, 16%, 17%, 18%, 19%, 20%, 25%, 30%, 35%, 40%, 45%, 50%, 55%, 60%, 65%, 70%, 75%, 80%, 85%, 90%, 95%, or any percentage in between these values.

[00181] The scanning magnet current may be controlled, in accordance with the treatment plan, to scan a depth-wise layer of the irradiation target. The layer is selected by appropriately positioning one or more energy degraders from the range compensator in the path of the particle beam and/or by setting an energy level of a

variable-energy particle accelerator. As the layer is scanned, the current sensor samples the current applied to the scanning magnet. The amount of magnet current may be recorded, e.g., stored in memory. If more than one magnet or magnet coil is used, the amount of magnet current may be stored along with the identity of the magnet or coil. In addition, the current may be correlated to coordinates within the irradiation target (e.g., Cartesian XYZ coordinates) and those coordinates may be stored in addition to, or instead of, the corresponding magnet current. As explained above, the current sensor may sample the magnet current and correlate the sampling time to the time at which an irradiation dose (e.g., pulse) is delivered.

[00182] In this regard, ion chamber 109 may detect the intensity of doses delivered to the irradiation target as that dose is delivered. The intensity of each dose is recorded (e.g., stored in memory) along with the location of each delivered dose. As noted, the location of each delivered dose may be stored by coordinates, magnet current, or using some other appropriate metric. As noted above, the dosimetry - the dose verification - may be synchronized with delivery of the dose and, thus, with the output of the accelerator (which corresponds to the RF frequency, as described above). Accordingly, in some implementations, each time a dose is delivered, the intensity of that dose is determined almost immediately and the location at which the dose is applied is determined almost immediately. This information may be stored in one or more tables (e.g., one table per layer or multiple tables per layer) or other appropriate computer storage construct.

[00183] In some implementations, the tables may be updated as additional doses are delivered. For example, a table may keep a running track of the amount of dose delivered at each location. So, if the beam dose is "X" grays, at a first scan pass, the table may record X grays for a location. At a second scan pass, the table may record 2X grays, and so forth until the target cumulative dose is reached.

[00184] In this regard, for each location, a processing device associated with the accelerator (e.g., the computer system that controls the particle therapy system) may compare the cumulative dose from a table, such as that described above, to the target cumulative dose. If the cumulative dose matches the target cumulative dose, treatment for that location (or layer) is deemed completed. If the cumulative dose

does not match the target cumulative dose, additional treatment is performed. For example, the layer or location is scanned again at the same locations, which are obtained from the table. The linear correlation between magnet current and beam movement produced by use of an air core magnet can facilitate repeated, and relatively accurate, repeated scanning at the same locations during multiple passes of the beam during scanning. The leaves of the adaptive aperture described herein may be moved slightly (e.g., a fraction of a millimeter, a millimeter, or multiple millimeters) between each scan of the same area, as described above.

[00185] Scanning may be repeated, at the same locations, any appropriate number of times until the target cumulative dose is reached at each location. In this regard, the entire layer may be re-scanned or only select portions of the layer may be re-scanned, dependent upon the target cumulative doses for the different locations on the layer. In some implementations, the intensity of the particle beam is not varied between scans. In other implementations, the intensity of the particle beam may be varied between scans, particularly if a small dose is required to top-off a cumulative dose to reach the target cumulative dose. The intensity of the dose may be increased or decreased, e.g., by altering the operation of the ion source (e.g., increasing the plasma ionization), altering the sweep of the RF frequency, or by any other appropriate methods.

[00186] As noted, scanning may be repeated for an entire layer or for only a portion of a layer. In some implementations, an entire layer, or a portion thereof, may be fully treated before treating another layer. That is, scanning may be repeated until the total cumulative dose is reached for each location on a layer before another layer is treated. In some implementations, each layer may be treated partially (e.g., scanned once) in sequence, and then re-scanned in sequence. In some implementations, several designated layers may be completely treated before other layers are treated. In some implementations, the entire target may be scanned once, followed by successive scans of the entire target until the appropriate total cumulative dose is delivered to each location. The adaptive aperture may be reconfigured for each scan or its configuration may remain static between scans.

[00187] During movement between layers, the beam may be turned-off. For example, during movement between layers, the ion source may be turned-off, thereby disrupting the output of the beam. During movement between layers, the RF sweep in the particle accelerator may be turned-off, thereby disrupting the extraction (and thus output) of the beam. During movement between layers, both the ion source and the circuitry that creates the RF sweep may be deactivated in some implementations. In some implementations, rather than turning-off the ion source and/or the RF sweep during movement between layers, the beam may be deflected to a beam-absorbing material using a kicker magnet (not shown) or the scanning magnet.

[00188] Different cross-sections of the irradiation target may be scanned according to different treatment plans. As described above, an energy degrader is used to control the scanning depth. In some implementations, the particle beam may be interrupted or redirected during configuration of the energy degrader. In other implementations, this need not be the case.

[00189] Described herein are examples of treating cross-sections of an irradiation target. These may be cross-sections that are roughly perpendicular to the direction of the particle beam. However, the concepts described herein are equally applicable to treating other portions of an irradiation target that are not cross-sections perpendicular to the direction of the particle beam. For example, an irradiation target may be segmented into spherical, cubical or other shaped volumes, and those volumes may be treated using the example processes, systems, and/or devices described herein.

[00190] Fig. 17 is a flowchart showing an example implementation of the scanning process described herein. Although process 200 of Fig. 17 is described in the context of the hardware described herein, process 200 may be performed using any appropriate hardware. The operations shown in process 200 may be performed in the same order as depicted or in a different order, where appropriate.

[00191] According to process 200, a treatment plan is stored (201). The treatment plan may be a treatment plan as described above. For example, the treatment plan may specify the type of scanning (e.g., uninterrupted raster scanning)

and the total cumulative dose of radiation to be delivered to each location in each layer of an irradiation target. In some implementations, the treatment plan may omit, e.g., the doses to be delivered for each scan at individual locations and their intensities, as well as the number of doses to be delivered to each location and the identity of the locations. In some implementations, the treatment plan may include, e.g., the doses to be delivered for each scan at individual locations and their intensities, as well as the number of doses to be delivered to each location and the identity of the locations.

[00192] An energy degrader may be set to select (202) a layer, and current may be applied to the magnet and controlled to move (203) the particle beam in accordance with a pattern set forth, e.g., in the treatment plan, to scan the layer. The current control may produce uninterrupted movement of the beam across at least part of the irradiation target to deliver doses of charged particles. An example of a pattern of beam movement 230 across a layer 233 of an irradiation target is shown in Fig. 18. As the beam moves, each pulse of the beam delivers a dose of radiation to the target. The dose has an intensity, which may be set beforehand in the accelerator or during scanning, and is delivered to specific positions. The exact positions at which dose is to be delivered need not be set beforehand, but rather may be arrived at by a combination of beam movement and pulse output.

[00193] For positions at which the dose is delivered, information is stored (204) (or otherwise recorded), which identifies a location and an amount of dose delivered at the location. This information is typically stored after the dose is delivered. As explained above, the information may be determined as close to delivery of the dose as possible, using the ion chamber to determine particle beam intensity (e.g., the amount of the dose) and the current sensor on the scanning magnet to determine the location at which the dose is delivered. As described above, in some implementations, in synchronism with the delivery, information identifying doses of the particle beam delivered to the irradiation target is stored along with at least one of: coordinates at which the doses were delivered or magnet currents at which the doses were delivered. As also described above, this

information may be stored in tables, which may be used to store the cumulative dose of radiation applied at positions on various layers of an irradiation target.

[00194] The entire layer may be scanned and information therefor recorded, as described above, or only part of the layer may be scanned and information therefor recorded. The adaptive aperture is mounted on an appropriate location of the scanning system in such a manner as to enable the adaptive aperture to track the particle beam's movement during scanning operations. At a point during scanning, the cumulative dose delivered at each position is compared to a target cumulative dose for that position. For example, this may be done after part of a layer containing that position is scanned, after the entire layer is scanned, after a set of layers are scanned, or after all layers in an irradiation target are scanned. It is determined (205) if the current cumulative dose matches the target cumulative dose at specific positions. If the current cumulative dose does match the target cumulative dose at specific positions, scanning is completed (207) for those positions. If the current cumulative dose does not match the target cumulative dose at specific positions, the scanning system is operated to compensate for deficiencies in the recorded (e.g., current cumulative) doses relative to corresponding target cumulative doses for those positions. For example, if the current cumulative dose does not match the target cumulative dose at specific positions, the current in the scanning magnet may be controlled in order to move (206) the beam so as to deliver additional dose to the specific positions.

[00195] As explained above, in some implementations, 100% of the dose may be applied during a single scan (e.g., a single delivery of particles) of a layer. In that case, more than one scan per layer may not be necessary. In other implementations, less than 100% of the dose may be applied during a single scan. In that case, more than one scan per layer will be necessary. To this end, according to the scanning process, for positions at which dose is applied, if the current cumulative dose at each position does not match the target cumulative dose at a corresponding position, the magnet current is controlled in order to move the beam so as to deliver additional dose to positions that require more dose. In other words, the layer may be re-scanned any appropriate number of times until the target

cumulative dose is reached for all positions of the layer. In some implementations, in one scan or in multiple scans, the actual delivered dose may exceed 100% of the target cumulative dose. What dose to deliver may be dictated by appropriate medical professionals.

[00196] As noted above, the layer may be re-scanned at any appropriate point, e.g., after part of the layer is completed with a current scan, after the entire layer is completed with the current scan, after a set of layers is completed with a scan, or after all layers are completed with a scan. During re-scanning, the process above is repeated until the target cumulative dose is reached for all, or some subset of, positions in the irradiation target. In some implementations, the intensity of the particle beam may need to be adjusted, e.g., for the last scan. For example, if the intensity is set at 25% of the target cumulative dose, but only 20% is delivered at each scan, then a fifth (and possibly sixth) dose will require a lower intensity than 25% in order to reach the target cumulative dose.

[00197] In some implementations, information from the TPP may be used to configure the adaptive aperture for an initial spot or series of spots. As described herein, configuring may include, for example, positioning carriages 713, 714, and/or 715, and/or one or more of leaves 735 and/or 736. In some implementations, the initial configuration may be made based on assumptions included in the treatment planning process including, but not limited to, the position of the irradiation target and patient during treatment. In some cases, however, the position of the irradiation target and patient during treatment may change relative to the assumptions included in the TPP. For example, the patient may move, breathing may impact target positions, and so forth. To account for these changes caused by movement, spot characteristics including, but not limited to, spot location, spot size, spot shape, spot weight, or spot orientation, may be correspondingly changed in order to implement the treatment. In some implementations, rather than modifying inputs to the TPP in order to make these changes via the TPP, the changes may be made online – e.g., during treatment while the accelerator is operational – without requiring any, or some, additional input from the TPP or the TPS.

[00198] For example, in some implementations, the position of the irradiation target may be tracked directly within the treatment room. For example, the treatment room may include one or more imaging systems to image and thereby track motion of the irradiation target. Examples of in-room imaging systems that may be used to track motion of an irradiation target are described in U.S. Patent Publication No. 2016/0067525, which was filed on September 9, 2014, and which is titled "Patient Positioning System". U.S. Patent Publication No. 2016/0067525 is incorporated by reference, at least for its disclosure of systems used to locate, and track, motion of an irradiation target in a treatment room.

[00199] In some implementations, the imaging system can be a three-dimensional (3D) imaging system. In some implementations, the imaging system can be a two-dimensional (2D) imaging system. In some implementations, the imaging system may be, or include, one or more of the following types of imaging systems: a CT system, an X-ray device to capture a radiograph image, a magnetic resonance imaging (MRI) device to capture an MRI image, a positron emission tomography (PET) device to capture a PET image, a SPECT/CT device (where SPECT is Single-photon emission computed tomography), and/or a video surface imaging device, or any combination of these imaging devices. In some implementations, images may be captured at different points in time so as to enable tracking of movement of the irradiation target. In some implementations, fiducials associated the patient and proximate to an irradiation target may be used to detect motion of the patient and/or irradiation target. The location of the fiducials relative to the target may be known, and movement of the fiducials may be detected and correlated to movement of the irradiation target. In some implementations, the irradiation target may be imaged directly at two different points in time.

[00200] In some implementations, movement may be tracked between application of individual spots on a single layer, between application of sets of spots (e.g., a predefined number of spots) on a single layer, between layers, or between some appropriate combination of the foregoing. Referring to example process 2500 of Fig. 43, an image of an irradiation target or patient is captured (2501) in order to determine the location of the irradiation target. The particle therapy system receives

(2502) information from a TPP/TPS that is used to produce (2503) a spot or set of spots on the irradiation target at appropriate locations based on the information. Thereafter, an image of an irradiation target or patient is captured (2504) in order to determine (2505) whether the irradiation target has moved since a prior image was captured. In some implementations, an X-ray system may be used to capture the image; however, any appropriate imaging system including, but not limited to, those described herein may be used. An irradiation target is deemed to have moved if the irradiation target is not in the same location, e.g., relative to some reference point, such as a fiducial, between images. If the irradiation target has not moved, one or more spots are produced (2503) using the most current treatment planning information. For example, the most current treatment planning information may be the information received (2502) from the TPP/TPS, or the most current treatment planning information may be the information as modified (2506) to account for prior movement. If the irradiation target has moved, the treatment planning information is modified (2506) to produce updated treatment planning information based on the movement of the irradiation target. Thereafter, spots are produced using the updated treatment planning information. The process may proceed as shown in Fig. 43 until the entire target is treated.

[00201] Thus, in some implementations, the TPP need not be re-run using new information, such as the newly-determined location of the irradiation target. Instead, in some implementations, the information from the TPP/TPS used to produce the spots is modified or changed (2506), e.g., in real-time, to account for the movement of the irradiation target. For example, the information may be changed while a particle accelerator, such as a synchrocyclotron, used to provide the treatment is online. Any appropriate information may be changed including, but not limited to, spot location, spot size, spot shape, spot weight, or spot orientation. For example, the locations of spots may be changed to account for movement of the irradiation target. In another example, the direction of application of particles to the irradiation target may be changed to account for movement of the irradiation target.

[00202] The processes described herein may be used with a single particle accelerator, and any two or more of the features thereof described herein may be

used with the single particle accelerator. The particle accelerator may be used in any type of medical or non-medical application. Examples of particle therapy systems that may be used are provided below. Notably, the concepts described herein may be used in other systems not specifically described.

[00203] Referring to Fig. 19, an example implementation of a charged particle radiation therapy system 400 includes a beam-producing particle accelerator 402 (e.g., the particle accelerator of Figs. 8 and 9) having a weight and size small enough to permit it to be mounted on a rotating gantry 404 with its output directed straight (that is, essentially directly) from the accelerator housing toward a patient 406. Particle accelerator 402 also includes a scanning system of a type described herein (e.g., in Figs. 10 to 18).

[00204] In some implementations, the steel gantry has two legs 408, 410 mounted for rotation on two respective bearings 412, 414 that lie on opposite sides of the patient. The accelerator is supported by a steel truss 416 that is long enough to span a treatment area 418 in which the patient lies (e.g., twice as long as a tall person, to permit the person to be rotated fully within the space with any desired target area of the patient remaining in the line of the beam) and is attached stably at both ends to the rotating legs of the gantry.

[00205] In some examples, the rotation of the gantry is limited to a range 420 of less than 360 degrees, e.g., about 180 degrees, to permit a floor 422 to extend from a wall of a vault 424 that houses the therapy system into the patient treatment area. The limited rotation range of the gantry also reduces the required thickness of some of the walls (which are not directly aligned with the beam, e.g., wall 430), which provide radiation shielding of people outside the treatment area. A range of 180 degrees of gantry rotation is enough to cover all treatment approach angles, but providing a larger range of travel can be useful. For example the range of rotation may be between 180 and 330 degrees and still provide clearance for the therapy floor space. In other implementations, rotation is not limited as described above.

[00206] The horizontal rotational axis 432 of the gantry is located nominally one meter above the floor where the patient and therapist interact with the therapy system. This floor is positioned about 3 meters above the bottom floor of the

therapy system shielded vault. The accelerator can swing under the raised floor for delivery of treatment beams from below the rotational axis. The patient couch moves and rotates in a substantially horizontal plane parallel to the rotational axis of the gantry. The couch can rotate through a range 434 of about 270 degrees in the horizontal plane with this configuration. This combination of gantry and patient rotational ranges and degrees of freedom allows the therapist to select virtually any approach angle for the beam. If needed, the patient can be placed on the couch in the opposite orientation and then all possible angles can be used.

[00207] In some implementations, the accelerator uses a synchrocyclotron configuration having a high magnetic field superconducting electromagnetic structure. Because the bend radius of a charged particle of a given kinetic energy is reduced in direct proportion to an increase in the magnetic field applied to it, the high magnetic field superconducting magnetic structure permits the accelerator to be made smaller and lighter. The synchrocyclotron uses a magnetic field that is uniform in rotation angle and falls off in strength with increasing radius. Such a field shape can be achieved regardless of the magnitude of the magnetic field, so in theory there is no upper limit to the magnetic field strength (and therefore the resulting particle energy at a fixed radius) that can be used in a synchrocyclotron.

[00208] The synchrocyclotron is supported on the gantry so that the beam is generated directly in line with the patient. The gantry permits rotation of the synchrocyclotron about a horizontal rotational axis that contains a point (isocenter 440) within, or near, the patient. The split truss that is parallel to the rotational axis, supports the synchrocyclotron on both sides.

[00209] Because the rotational range of the gantry is limited in some example implementations, a patient support area can be accommodated in a wide area around the isocenter. Because the floor can be extended broadly around the isocenter, a patient support table can be positioned to move relative to and to rotate about a vertical axis 442 through the isocenter so that, by a combination of gantry rotation and table motion and rotation, any angle of beam direction into any part of the patient can be achieved. In some implementations, the two gantry arms are

separated by more than twice the height of a tall patient, allowing the couch with patient to rotate and translate in a horizontal plane above the raised floor.

[00210] Limiting the gantry rotation angle allows for a reduction in the thickness of at least one of the walls surrounding the treatment room. Thick walls, typically constructed of concrete, provide radiation protection to individuals outside the treatment room. A wall downstream of a stopping proton beam may be about twice as thick as a wall at the opposite end of the room to provide an equivalent level of protection. Limiting the range of gantry rotation enables the treatment room to be sited below earth grade on three sides, while allowing an occupied area adjacent to the thinnest wall reducing the cost of constructing the treatment room.

[00211] In the example implementation shown in Fig. 19, the superconducting synchrocyclotron 402 operates with a peak magnetic field in a pole gap of the synchrocyclotron of 8.8 Tesla. The synchrocyclotron produces a beam of protons having an energy of 250 MeV. In some implementations, the synchrocyclotron is a variable-energy machine, and is capable of outputting proton beams having different energies. In some implementations, the synchrocyclotron may produce a beam having a fixed energy. In some implementations the field strength could be in the range of 4T to 20T and the proton energy could be in the range of 100 to 300 MeV.

[00212] The radiation therapy system described in this example is used for proton radiation therapy, but the same principles and details can be applied in analogous systems for use in heavy ion (ion) treatment systems.

[00213] As shown in Figs. 8, 9, 20, 21, and 22, an example synchrocyclotron 10 (e.g., 402 in Fig. 19) includes a magnet system 122 that contains a particle source 190, a radiofrequency drive system 191, and a beam extraction system. In this example, the magnetic field established by the magnet system has a shape appropriate to maintain focus of a contained proton beam using a combination of a split pair of annular superconducting coils 140, 142 and a pair of shaped ferromagnetic (e.g., low carbon steel) pole faces 144, 146.

[00214] The two superconducting magnet coils are centered on a common axis and are spaced apart along the axis. The coils may be formed by of Nb₃Sn-based superconducting strands (that initially comprise a niobium-tin core surrounded by a

copper sheath) deployed in a twisted cable-in-channel conductor geometry. After individual strands are cabled together, they are heated to cause a reaction that forms the final (brittle) superconducting material of the wire. After the material has been reacted, the wires are soldered into the copper channel and covered with insulation. The copper channel containing the wires is then wound in a coil having a rectangular cross-section. The wound coil is then vacuum impregnated with an epoxy compound. The finished coils are mounted on an annular stainless steel reverse bobbin. Heater blankets may be placed at intervals in the layers of the windings to protect the assembly in the event of a magnet quench.

[00215] The entire coil can then be covered with copper sheets to provide thermal conductivity and mechanical stability and then contained in an additional layer of epoxy. The precompression of the coil can be provided by heating the stainless steel reverse bobbin and fitting the coils within the reverse bobbin. The reverse bobbin inner diameter is chosen so that when the entire mass is cooled to 4 K, the reverse bobbin stays in contact with the coil and provides some compression. Heating the stainless steel reverse bobbin to approximately 50 degrees C and fitting coils at a temperature of 100 degrees Kelvin can achieve this.

[00216] The geometry of the coil is maintained by mounting the coils in a "reverse" rectangular bobbin to exert a restorative force that works against the distorting force produced when the coils are energized. As shown in Fig. 21, in some implementations, coil position is maintained relative to corresponding magnet pole pieces and the cryostat using a set of warm-to-cold support straps 402, 404, 406. Supporting the cold mass with thin straps reduces the heat leakage imparted to the cold mass by the rigid support system. The straps are arranged to withstand the varying gravitational force on the coil as the magnet rotates on board the gantry. They withstand the combined effects of gravity and the large de-centering force realized by the coil when it is perturbed from a perfectly symmetric position relative to the magnet yoke. Additionally, the links act to reduce dynamic forces imparted on the coil as the gantry accelerates and decelerates when its position is changed. Each warm-to-cold support may include one S2 fiberglass link and one carbon fiber link. The carbon fiber link is supported across pins between the warm yoke and an

intermediate temperature (50 - 70 K), and the S2 fiberglass link 408 is supported across the intermediate temperature pin and a pin attached to the cold mass. Each pin may be made of high strength stainless steel.

[00217] Referring to Fig. 8, the field strength profile as a function of radius is determined largely by choice of coil geometry and pole face shape. The pole faces 144, 146 of the permeable yoke material can be contoured to fine tune the shape of the magnetic field to ensure that the particle beam remains focused during acceleration.

[00218] The superconducting coils are maintained at temperatures near absolute zero (e.g., about 4 degrees Kelvin) by enclosing the coil assembly (the coils and the bobbin) inside an evacuated annular aluminum or stainless steel cryostatic chamber 170 (the cryostat) that provides a free space around the coil structure, except at a limited set of support points 171, 173. In an alternate version (e.g., Fig. 9) the outer wall of the cryostat may be made of low carbon steel to provide an additional return flux path for the magnetic field.

[00219] In some implementations, the temperature near absolute zero is achieved and maintained using one single-stage Gifford-McMahon cryo-cooler and three two-stage Gifford McMahon cryo-coolers. Each two stage cryo-cooler has a second stage cold end attached to a condenser that recondenses Helium vapor into liquid Helium. In some implementations, the temperature near absolute zero is achieved and maintained using a cooling channel (not shown) containing liquid helium, which is formed inside a superconducting coil support structure (e.g., the reverse bobbin 156), and which contains a thermal connection between the liquid helium in the channel and the corresponding superconducting coil.

[00220] In some implementations, the coil assembly and cryostatic chambers are mounted within and fully enclosed by two halves 181, 183 of a pillbox-shaped magnet yoke 100. The yoke 100 provides a path for the return magnetic field flux 184 and magnetically shields the volume 186 between the pole faces 144, 146 to prevent external magnetic influences from perturbing the shape of the magnetic field within that volume. The yoke also serves to decrease the stray magnetic field in the vicinity of the accelerator. In other implementations, the coil assembly and cryostatic

chambers are mounted within and fully enclosed by a non-magnetic enclosure, and the path for return magnetic field flux is implemented using an active return system, an example of which is described above.

[00221] As shown in Figs. 8 and 23, the synchrocyclotron includes a particle source 190 of a Penning ion gauge geometry located near the geometric center 192 of the magnet structure. The particle source may be as described below.

[00222] Particle source 190 is fed from a supply 399 of hydrogen through a gas line 393 and tube 394 that delivers gaseous hydrogen. Electric cables 294 carry an electric current from a current source to stimulate electron discharge from cathodes 392, 390 that are aligned with the magnetic field. Particle source may be interrupted at the acceleration plane, e.g., the entirety of the particle source may be removed, thereby exposing the plasma column.

[00223] In this example, the discharged electrons ionize the gas exiting through a small hole from tube 394 to create a supply of positive ions (protons) for acceleration by one semicircular (dee-shaped) radio-frequency plate that spans half of the space enclosed by the magnet structure and one dummy dee plate. In the case of an interrupted particle source, all (or a substantial part, e.g., a majority) of the tube containing plasma is removed at the acceleration region.

[00224] As shown in Fig. 24, the dee plate 500 is a hollow metal structure that has two semicircular surfaces 503, 505 that enclose a space 507 in which the protons are accelerated during half of their rotation around the space enclosed by the magnet structure. A duct 509 opening into the space 507 extends through the enclosure (e.g., the yoke or pole piece(s)) to an external location from which a vacuum pump can be attached to evacuate the space 507 and the rest of the space within a vacuum chamber in which the acceleration takes place. The dummy dee 502 comprises a rectangular metal ring that is spaced near to the exposed rim of the dee plate. The dummy dee is grounded to the vacuum chamber and magnet yoke. The dee plate 500 is driven by a radio-frequency signal that is applied at the end of a radio-frequency transmission line to impart an electric field in the space 507. The radio frequency electric field is made to vary in time as the accelerated particle beam increases in distance from the geometric center.

[00225] For the beam emerging from the centrally located particle source to clear the particle source structure as it begins to spiral outward, a large voltage difference may be applied across the radio frequency plates. 20,000 Volts is applied across the radio frequency plates. In some versions from 8,000 to 20,000 Volts may be applied across the radio frequency plates. To reduce the power required to drive this large voltage, the magnet structure is arranged to reduce the capacitance between the radio frequency plates and ground. This may be done by forming holes with sufficient clearance from the radio frequency structures through the outer yoke and the cryostat housing and making sufficient space between the magnet pole faces.

[00226] The high voltage alternating potential that drives the dee plate has a frequency that is swept downward during the accelerating cycle to account for the increasing relativistic mass of the protons and the decreasing magnetic field. The dummy dee does not require a hollow semi-cylindrical structure as it is at ground potential along with the vacuum chamber walls. Other plate arrangements could be used such as more than one pair of accelerating electrodes driven with different electrical phases or multiples of the fundamental frequency. The RF structure can be tuned to keep the Q high during the required frequency sweep by using, for example, a rotating capacitor having intermeshing rotating and stationary blades. During each meshing of the blades, the capacitance increases, thus lowering the resonant frequency of the RF structure. The blades can be shaped to create a precise frequency sweep required. A drive motor for the rotating condenser can be phase locked to the RF generator for precise control. One bunch of particles may be accelerated during each meshing of the blades of the rotating condenser.

[00227] The vacuum chamber in which the acceleration occurs is a generally cylindrical container that is thinner in the center and thicker at the rim. The vacuum chamber encloses the RF plates and the particle source and is evacuated by a vacuum pump. Maintaining a high vacuum reduces the chances that accelerating ions are not lost to collisions with gas molecules and enables the RF voltage to be kept at a higher level without arcing to ground.

[00228] Protons (or other ions) traverse a generally spiral orbital path beginning at the particle source. In half of each loop of the spiral path, the protons gain energy as they pass through the RF electric field. As the protons gain energy, the radius of the central orbit of each successive loop of their spiral path is larger than the prior loop until the loop radius reaches the maximum radius of the pole face. At that location a magnetic and electric field perturbation directs protons into an area where the magnetic field rapidly decreases, and the protons depart the area of the high magnetic field and are directed through an evacuated tube, referred to herein as the extraction channel, to exit the synchrocyclotron. A magnetic regenerator may be used to change the magnetic field perturbation to direct the protons. The protons exiting will tend to disperse as they enter an area of markedly decreased magnetic field that exists in the room around the synchrocyclotron. Beam shaping elements 607, 609 in the extraction channel 138 (Fig. 21) redirect the protons so that they stay in a straight beam of limited spatial extent.

[00229] As the beam exits the extraction channel it is passed through a beam formation system 525 (Fig. 21), which may include a scanning system of the type described herein. Beam formation system 525 may be used in conjunction with an inner gantry that controls application of the beam.

[00230] Stray magnetic fields exiting from the synchrocyclotron may be limited by both a magnet yoke (which also serves as a shield) and a separate magnetic shield 514 (e.g., Fig. 8). The separate magnetic shield includes of a layer 517 of ferromagnetic material (e.g., steel or iron) that encloses the pillbox yoke, separated by a space 516. This configuration that includes a sandwich of a yoke, a space, and a shield achieves adequate shielding for a given leakage magnetic field at lower weight. As described above, in some implementations, an active return system may be used in place of, or to augment, the operation of the magnetic yoke and shield.

[00231] Referring to Fig. 19, the gantry allows the synchrocyclotron to be rotated about a horizontal rotational axis 432. The truss structure 416 has two generally parallel spans 480, 482. The synchrocyclotron is cradled between the spans about midway between the legs. The gantry is balanced for rotation about the

bearings using counterweights 622, 624 mounted on ends of the legs opposite the truss.

[00232] The gantry is driven to rotate by an electric motor mounted to one or both of the gantry legs and connected to the bearing housings by drive gears. The rotational position of the gantry is derived from signals provided by shaft angle encoders incorporated into the gantry drive motors and the drive gears.

[00233] At the location at which the ion beam exits the synchrocyclotron, the beam formation system 525 acts on the ion beam to give it properties suitable for patient treatment. For example, the beam may be spread and its depth of penetration varied to provide uniform radiation across a given target volume. The beam formation system may include active scanning elements as described herein.

[00234] All of the active systems of the synchrocyclotron (the current driven superconducting coils, the RF-driven plates, the vacuum pumps for the vacuum acceleration chamber and for the superconducting coil cooling chamber, the current driven particle source, the hydrogen gas source, and the RF plate coolers, for example), may be controlled by appropriate synchrocyclotron control electronics (not shown), which may include, e.g., one or more processing devices executing instructions from non-transitory memory to effect control.

[00235] The particle accelerator used in the example particle therapy systems and example scanning systems described herein may be a variable-energy particle accelerator, an example of which is described below

[00236] The energy of an extracted particle beam (the particle beam output from the accelerator) can affect the use of the particle beam during treatment. In some machines, the energy of the particle beam (or particles in the particle beam) does not increase after extraction. However, the energy may be reduced based on treatment needs after the extraction and before the treatment. Referring to Fig. 25, an example treatment system 910 includes an accelerator 912, e.g., a synchrocyclotron, from which a particle (e.g., proton) beam 914 having a variable energy is extracted to irradiate a target volume 924 of a body 922. Optionally, one or more additional devices, such as a scanning unit 916 or a scattering unit 916, one or more monitoring units 918, and an energy degrader 920, are placed along the

irradiation direction 928. The devices intercept the cross-section of the extracted beam 914 and alter one or more properties of the extracted beam for the treatment. The adaptive aperture may be located down-beam of the energy degrader, between the energy degrader and the patient.

[00237] A target volume to be irradiated (an irradiation target) by a particle beam for treatment typically has a three-dimensional configuration. In some examples, to carry-out the treatment, the target volume is divided into layers along the irradiation direction of the particle beam so that the irradiation can be done on a layer-by-layer basis. For certain types of particles, such as protons, the penetration depth (or which layer the beam reaches) within the target volume is largely determined by the energy of the particle beam. A particle beam of a given energy does not reach substantially beyond a corresponding penetration depth for that energy. To move the beam irradiation from one layer to another layer of the target volume, the energy of the particle beam is changed.

[00238] In the example shown in Fig. 25, the target volume 924 is divided into nine layers 926a-926i along the irradiation direction 928. In an example process, the irradiation starts from the deepest layer 926i, one layer at a time, gradually to the shallower layers and finishes with the shallowest layer 926a. Before application to the body 922, the energy of the particle beam 914 is controlled to be at a level to allow the particle beam to stop at a desired layer, e.g., the layer 926d, without substantially penetrating further into the body or the target volume, e.g., the layers 926e-926i or deeper into the body. In some examples, the desired energy of the particle beam 914 decreases as the treatment layer becomes shallower relative to the particle acceleration. In some examples, the beam energy difference for treating adjacent layers of the target volume 924 is about 3 MeV to about 100 MeV, e.g., about 10 MeV to about 80 MeV, although other differences may also be possible, depending on, e.g., the thickness of the layers and the properties of the beam.

[00239] The energy variation for treating different layers of the target volume 924 can be performed at the accelerator 912 (e.g., the accelerator can vary the energy) so that, in some implementations, no additional energy variation is required after the particle beam is extracted from the accelerator 912. So, the optional

energy degrader 920 in the treatment system 10 may be eliminated from the system. In some implementations, the accelerator 912 can output particle beams having an energy that varies between about 100 MeV and about 300 MeV, e.g., between about 115 MeV and about 250 MeV. The variation can be continuous or non-continuous, e.g., one step at a time. In some implementations, the variation, continuous or non-continuous, can take place at a relatively high rate, e.g., up to about 50 MeV per second or up to about 20 MeV per second. Non-continuous variation can take place one step at a time with a step size of about 10 MeV to about 90 MeV.

[00240] When irradiation is complete in one layer, the accelerator 912 can vary the energy of the particle beam for irradiating a next layer, e.g., within several seconds or within less than one second. In some implementations, the treatment of the target volume 924 can be continued without substantial interruption or even without any interruption. In some situations, the step size of the non-continuous energy variation is selected to correspond to the energy difference needed for irradiating two adjacent layers of the target volume 924. For example, the step size can be the same as, or a fraction of, the energy difference.

[00241] In some implementations, the accelerator 912 and the degrader 920 collectively vary the energy of the beam 914. For example, the accelerator 912 provides a coarse adjustment and the degrader 920 provides a fine adjustment or vice versa. In this example, the accelerator 912 can output the particle beam that varies energy with a variation step of about 10-80 MeV, and the degrader 920 adjusts (e.g., reduces) the energy of the beam at a variation step of about 2-10 MeV.

[00242] The reduced use (or absence) of the energy degrader, such as a range modulator, may help to maintain properties and quality of the output beam from the accelerator, e.g., beam intensity. The control of the particle beam can be performed at the accelerator. Side effects, e.g., from neutrons generated when the particle beam passes the degrader 920 can be reduced or eliminated.

[00243] The energy of the particle beam 914 may be adjusted to treat another target volume 930 in another body or body part 922' after completing treatment in target volume 924. The target volumes 924, 930 may be in the same body (or patient), or in different patients. It is possible that the depth D of the target volume

930 from a surface of body 922' is different from that of the target volume 924.

Although some energy adjustment may be performed by the degrader 920, the degrader 912 may only reduce the beam energy and not increase the beam energy.

[00244] In this regard, in some cases, the beam energy required for treating target volume 930 is greater than the beam energy required to treat target volume 924. In such cases, the accelerator 912 may increase the output beam energy after treating the target volume 924 and before treating the target volume 930. In other cases, the beam energy required for treating target volume 930 is less than the beam energy required to treat target volume 924. Although the degrader 920 can reduce the energy, the accelerator 912 can be adjusted to output a lower beam energy to reduce or eliminate the use of the degrader 920. The division of the target volumes 924, 930 into layers can be different or the same. The target volume 930 can be treated similarly on a layer by layer basis to the treatment of the target volume 924.

[00245] The treatment of the different target volumes 924, 930 on the same patient may be substantially continuous, e.g., with the stop time between the two volumes being no longer than about 30 minutes or less, e.g., 25 minutes or less, 20 minutes or less, 15 minutes or less, 10 minutes or less, 5 minutes or less, or 1 minute or less. As explained herein, the accelerator 912 can be mounted on a movable gantry and the movement of the gantry can move the accelerator to aim at different target volumes. In some situations, the accelerator 912 can complete the energy adjustment of the output beam 914 during the time the treatment system makes adjustment (such as moving the gantry) after completing the treatment of the target volume 924 and before starting treating the target volume 930. After the alignment of the accelerator and the target volume 930, the treatment can begin with the adjusted, desired beam energy. Beam energy adjustment for different patients can also be completed relatively efficiently. In some examples, all adjustments, including increasing/reducing beam energy and/or moving the gantry are done within about 30 minutes, e.g., within about 25 minutes, within about 20 minutes, within about 15 minutes, within about 10 minutes or within about 5 minutes.

[00246] In the same layer of a target volume, an irradiation dose may be applied by moving the beam across the two-dimensional surface of the layer (which is sometimes called scanning beam) using a scanning unit 916. Alternatively, the layer can be irradiated by passing the extracted beam through one or more scatterers of the scattering unit 16 (which is sometimes called scattering beam).

[00247] Beam properties, such as energy and intensity, can be selected before a treatment or can be adjusted during the treatment by controlling the accelerator 912 and/or other devices, such as the scanning unit/scatterer(s) 916, the degrader 920, and others not shown in the figures. In example implementations, system 910 includes a controller 932, such as a computer, in communication with one or more devices in the system. Control can be based on results of the monitoring performed by the one or more monitors 918, e.g., monitoring of the beam intensity, dose, beam location in the target volume, etc. Although the monitors 918 are shown to be between the device 916 and the degrader 920, one or more monitors can be placed at other appropriate locations along the beam irradiation path. Controller 932 can also store a treatment plan for one or more target volumes (for the same patient and/or different patients). The treatment plan can be determined before the treatment starts and can include parameters, such as the shape of the target volume, the number of irradiation layers, the irradiation dose for each layer, the number of times each layer is irradiated, etc. The adjustment of a beam property within the system 910 can be performed based on the treatment plan. Additional adjustment can be made during the treatment, e.g., when deviation from the treatment plan is detected.

[00248] In some implementations, the accelerator 912 is configured to vary the energy of the output particle beam by varying the magnetic field in which the particle beam is accelerated. In an example implementation, one or more sets of coils receives variable electrical current to produce a variable magnetic field in the cavity. In some examples, one set of coils receives a fixed electrical current, while one or more other sets of coils receives a variable current so that the total current received by the coil sets varies. In some implementations, all sets of coils are superconducting. In other implementations, some sets of coils, such as the set for

the fixed electrical current, are superconducting, while other sets of coils, such as the one or more sets for the variable current, are non-superconducting. In some examples, all sets of coils are non-superconducting.

[00249] Generally, the magnitude of the magnetic field is scalable with the magnitude of the electrical current. Adjusting the total electric current of the coils in a predetermined range can generate a magnetic field that varies in a corresponding, predetermined range. In some examples, a continuous adjustment of the electrical current can lead to a continuous variation of the magnetic field and a continuous variation of the output beam energy. Alternatively, when the electrical current applied to the coils is adjusted in a non-continuous, step-wise manner, the magnetic field and the output beam energy also varies accordingly in a non-continuous (step-wise) manner. The scaling of the magnetic field to the current can allow the variation of the beam energy to be carried out relatively precisely, although sometimes minor adjustment other than the input current may be performed.

[00250] In some implementations, to output particle beams having a variable energy, the accelerator 912 is configured to apply RF voltages that sweep over different ranges of frequencies, with each range corresponding to a different output beam energy. For example, if the accelerator 912 is configured to produce three different output beam energies, the RF voltage is capable of sweeping over three different ranges of frequencies. In another example, corresponding to continuous beam energy variations, the RF voltage sweeps over frequency ranges that continuously change. The different frequency ranges may have different lower frequency and/or upper frequency boundaries.

[00251] The extraction channel may be configured to accommodate the range of different energies produced by the variable-energy particle accelerator. For example the extraction channel may be large enough to support the highest and lowest energies produced by the particle accelerator. That is, the extraction channel may be sized or otherwise configured to receive and to transmit particles within that range of energies. Particle beams having different energies can be extracted from the accelerator 912 without altering the features of the regenerator that is used for extracting particle beams having a single energy. In other implementations, to

accommodate the variable particle energy, the regenerator can be moved to disturb (e.g., change) different particle orbits in the manner described above and/or iron rods (magnetic shims) can be added or removed to change the magnetic field bump provided by the regenerator. More specifically, different particle energies will typically be at different particle orbits within the cavity. By moving the regenerator, it is possible to intercept a particle orbit at a specified energy and thereby provide the correct perturbation of that orbit so that particles at the specified energy reach the extraction channel. In some implementations, movement of the regenerator (and/or addition/removal of magnetic shims) is performed in real-time to match real-time changes in the particle beam energy output by the accelerator. In other implementations, particle energy is adjusted on a per-treatment basis, and movement of the regenerator (and/or addition/removal of magnetic shims) is performed in advance of the treatment. In either case, movement of the regenerator (and/or addition/removal of magnetic shims) may be computer controlled. For example, a computer may control one or more motors that effect movement of the regenerator and/or magnetic shims.

[00252] In some implementations, the regenerator is implemented using one or more magnetic shims that are controllable to move to the appropriate location(s).

[00253] As an example, table 1 shows three example energy levels at which example accelerator 912 can output particle beams. The corresponding parameters for producing the three energy levels are also listed. In this regard, the magnet current refers to the total electrical current applied to the one or more coil sets in the accelerator 912; the maximum and minimum frequencies define the ranges in which the RF voltage sweeps; and "r" is the radial distance of a location to a center of the cavity in which the particles are accelerated.

Beam Energy (MeV)	Magnet Current (Amps)	Maximum Frequency (MHz)	Minimum Frequency (MHz)	Magnetic Field at r=0 mm (Tesla)	Magnetic Field at r=298 mm (Tesla)
250	1990	132	99	8.7	8.2
235	1920	128	97	8.4	8.0
211	1760	120	93	7.9	7.5

Table 1. Examples of beam energies and respective parameters.

[00254] The accelerator may be a synchrocyclotron and the particles may be protons. The particles may be output as pulsed beams. The energy of the beam output from the particle accelerator can be varied during the treatment of one target volume in a patient, or between treatments of different target volumes of the same patient or different patients. In some implementations, settings of the accelerator are changed to vary the beam energy when no beam (or particles) is output from the accelerator. The energy variation can be continuous or non-continuous over a desired range.

[00255] Referring to the example shown in Fig. 8, the particle accelerator, which may be a variable-energy particle accelerator like accelerator 912 described above, may be configured to output particle beams that have a variable energy. The range of the variable energy can have an upper boundary that is about 200 MeV to about 300 MeV or higher, e.g., 200 MeV, about 205 MeV, about 210 MeV, about 215 MeV, about 220 MeV, about 225 MeV, about 230 MeV, about 235 MeV, about 240 MeV, about 245 MeV, about 250 MeV, about 255 MeV, about 260 MeV, about 265 MeV, about 270 MeV, about 275 MeV, about 280 MeV, about 285 MeV, about 290 MeV, about 295 MeV, or about 300 MeV or higher. The range can also have a lower boundary that is about 100 MeV or lower to about 200 MeV, e.g., about 100 MeV or lower, about 105 MeV, about 110 MeV, about 115 MeV, about 120 MeV, about 125 MeV, about 130 MeV, about 135 MeV, about 140 MeV, about 145 MeV, about 150 MeV, about 155 MeV, about 160 MeV, about 165 MeV, about 170 MeV, about 175 MeV, about 180 MeV, about 185 MeV, about 190 MeV, about 195 MeV, about 200 MeV.

[00256] In some examples, the variation is non-continuous and the variation step can have a size of about 10 MeV or lower, about 15 MeV, about 20 MeV, about 25 MeV, about 30 MeV, about 35 MeV, about 40 MeV, about 45 MeV, about 50 MeV, about 55 MeV, about 60 MeV, about 65 MeV, about 70 MeV, about 75 MeV, or about 80 MeV or higher. Varying the energy by one step size can take no more than 30 minutes, e.g., about 25 minutes or less, about 20 minutes or less, about 15 minutes or less, about 10 minutes or less, about 5 minutes or less, about 1 minute or less, or about 30 seconds or less. In other examples, the variation is continuous and

the accelerator can adjust the energy of the particle beam at a relatively high rate, e.g., up to about 50 MeV per second, up to about 45 MeV per second, up to about 40 MeV per second, up to about 35 MeV per second, up to about 30 MeV per second, up to about 25 MeV per second, up to about 20 MeV per second, up to about 15 MeV per second, or up to about 10 MeV per second. The accelerator can be configured to adjust the particle energy both continuously and non-continuously. For example, a combination of the continuous and non-continuous variation can be used in a treatment of one target volume or in treatments of different target volumes. Flexible treatment planning and flexible treatment can be achieved.

[00257] A particle accelerator that outputs a particle beam having a variable energy can provide accuracy in irradiation treatment and reduce the number of additional devices (other than the accelerator) used for the treatment. For example, the use of degraders for changing the energy of an output particle beam may be reduced or eliminated for all or part of the treatment. The properties of the particle beam, such as intensity, focus, etc. can be controlled at the particle accelerator and the particle beam can reach the target volume without substantial disturbance from the additional devices. The relatively high variation rate of the beam energy can reduce treatment time and allow for efficient use of the treatment system.

[00258] In some implementations, the accelerator, such as the synchrocyclotron of Fig. 8, accelerates particles or particle beams to variable energy levels by varying the magnetic field in the accelerator, which can be achieved by varying the electrical current applied to coils for generating the magnetic field. As explained above, an example synchrocyclotron (e.g., the synchrocyclotron of Fig. 8) includes a magnet system that contains a particle source, a radiofrequency drive system, and a beam extraction system. Fig. 26 shows an example of a magnet system that may be used in a variable-energy accelerator. In this example implementation, the magnetic field established by the magnet system 1012 can vary by about 5% to about 35% of a maximum value of the magnetic field that two sets of coils 40a and 40b, and 42a and 42b are capable of generating. The magnetic field established by the magnet system has a shape appropriate to maintain focus of a contained proton beam using a combination of the two sets of coils and a pair of

shaped ferromagnetic (e.g., low carbon steel) structures, examples of which are provided above.

[00259] Each set of coils may be a split pair of annular coils to receive electrical current. In some situations, both sets of coils are superconducting. In other situations, only one set of the coils is superconducting and the other set is non-superconducting or normal conducting (also discussed further below). It is also possible that both sets of coils are non-superconducting. Suitable superconducting materials for use in the coils include niobium-3 tin (Nb₃Sn) and/or niobium-titanium. Other example conducting materials can include copper. Examples of the coil set constructions are described further below.

[00260] The two sets of coils can be electrically connected serially or in parallel. In some implementations, the total electrical current received by the two sets of coils can include about 2 million ampere turns to about 10 million ampere turns, e.g., about 2.5 to about 7.5 million ampere turns or about 3.75 million ampere turns to about 5 million ampere turns. In some examples, one set of coils is configured to receive a fixed (or constant) portion of the total variable electrical current, while the other set of coils is configured to receive a variable portion of the total electrical current. The total electrical current of the two coil sets varies with the variation of the current in one coil set. In other situations, the electrical current applied to both sets of coils can vary. The variable total current in the two sets of coils can generate a magnetic field having a variable magnitude, which in turn varies the acceleration pathways of the particles and produces particles having variable energies.

[00261] Generally, the magnitude of the magnetic field generated by the coil(s) is scalable to the magnitude of the total electrical current applied to the coil(s). Based on the scalability, in some implementations, linear variation of the magnetic field strength can be achieved by linearly changing the total current of the coil sets. The total current can be adjusted at a relatively high rate that leads to a relatively high-rate adjustment of the magnetic field and the beam energy.

[00262] In the example reflected in Table 1 above, the ratio between values of the current and the magnetic field at the geometric center of the coil rings is:

1990:8.7 (approximately 228.7:1); 1920:8.4 (approximately 228.6:1); 1760:7.9 (approximately 222.8:1). Accordingly, adjusting the magnitude of the total current applied to a superconducting coil(s) can proportionally (based on the ratio) adjust the magnitude of the magnetic field.

[00263] The scalability of the magnetic field to the total electrical current in the example of Table 1 is also shown in the plot of Fig. 27, where BZ is the magnetic field along the Z-dimension; and R is the radial distance measured from a geometric center of the coil rings along a direction perpendicular to the Z-dimension. The magnetic field has the highest value at the geometric center, and decreases as the distance R increases. The curves 1035, 1037 represent the magnetic field generated by the same coil sets receiving different total electrical current: 1760 Amperes and 1990 Amperes, respectively. The corresponding energies of the extracted particles are 211 MeV and 250 MeV, respectively. The two curves 1035, 1037 have substantially the same shape and the different parts of the curves 1035, 1037 are substantially parallel. As a result, either the curve 1035 or the curve 1037 can be linearly shifted to substantially match the other curve, indicating that the magnetic field is scalable to the total electrical current applied to the coil sets.

[00264] In some implementations, the scalability of the magnetic field to the total electrical current may not be perfect. For example, the ratio between the magnetic field and the current calculated based on the example shown in table 1 is not constant. Also, as shown in Fig. 27, the linear shift of one curve may not perfectly match the other curve. In some implementations, the total current is applied to the coil sets under the assumption of perfect scalability. The target magnetic field (under the assumption of perfect scalability) can be generated by additionally altering the features, e.g., geometry, of the coils to counteract the imperfection in the scalability. As one example, ferromagnetic (e.g., iron) rods (magnetic shims) can be inserted or removed from one or both of the magnetic structures (e.g., yokes, pole pieces, and the like). The features of the coils can be altered at a relatively high rate so that the rate of the magnetic field adjustment is not substantially affected as compared to the situation in which the scalability is perfect and only the electrical current needs to be adjusted. In the example of iron rods, the

rods can be added or removed at the time scale of seconds or minutes, e.g., within 5 minutes, within 1 minute, less than 30 seconds, or less than 1 second.

[00265] In some implementations, settings of the accelerator, such as the current applied to the coil sets, can be chosen based on the substantial scalability of the magnetic field to the total electrical current in the coil sets.

[00266] Generally, to produce the total current that varies within a desired range, any appropriate combination of current applied to the two coil sets can be used. In an example, the coil set 42a, 42b can be configured to receive a fixed electrical current corresponding to a lower boundary of a desired range of the magnetic field. In the example shown in table 1, the fixed electrical current is 1760 Amperes. In addition, the coil set 40a, 40b can be configured to receive a variable electrical current having an upper boundary corresponding to a difference between an upper boundary and a lower boundary of the desired range of the magnetic field. In the example shown in table 1, the coil set 40a, 40b is configured to receive electrical current that varies between 0 Ampere and 230 Amperes.

[00267] In another example, the coil set 42a, 42b can be configured to receive a fixed electrical current corresponding to an upper boundary of a desired range of the magnetic field. In the example shown in table 1, the fixed current is 1990 Amperes. In addition, the coil set 40a, 40b can be configured to receive a variable electrical current having an upper boundary corresponding to a difference between a lower boundary and an upper boundary of the desired range of the magnetic field. In the example shown in table 1, the coil set 40a, 40b is configured to receive electrical current that varies between -230 Ampere and 0 Ampere.

[00268] The total variable magnetic field generated by the variable total current for accelerating the particles can have a maximum magnitude greater than 4 Tesla, e.g., greater than 5 Tesla, greater than 6 Tesla, greater than 7 Tesla, greater than 8 Tesla, greater than 9 Tesla, or greater than 10 Tesla, and up to about 20 Tesla or higher, e.g., up to about 18 Tesla, up to about 15 Tesla, or up to about 12 Tesla. In some implementations, variation of the total current in the coil sets can vary the magnetic field by about 0.2 Tesla to about 4.2 Tesla or more, e.g., about 0.2 Tesla to about 1.4 Tesla or about 0.6 Tesla to about 4.2 Tesla. In some situations, the

amount of variation of the magnetic field can be proportional to the maximum magnitude.

[00269] Fig. 28 shows an example RF structure for sweeping the voltage on the dee plate 500 over an RF frequency range for each energy level of the particle beam, and for varying the frequency range when the particle beam energy is varied. The semicircular surfaces 503, 505 of the dee plate 500 are connected to an inner conductor 1300 and housed in an outer conductor 1302. The high voltage is applied to the dee plate 500 from a power source (not shown, e.g., an oscillating voltage input) through a power coupling device 1304 that couples the power source to the inner conductor. In some implementations, the coupling device 1304 is positioned on the inner conductor 1300 to provide power transfer from the power source to the dee plate 500. In addition, the dee plate 500 is coupled to variable reactive elements 1306, 1308 to perform the RF frequency sweep for each particle energy level, and to change the RF frequency range for different particle energy levels.

[00270] The variable reactive element 1306 can be a rotating capacitor that has multiple blades 1310 rotatable by a motor (not shown). By meshing or unmeshing the blades 1310 during each cycle of RF sweeping, the capacitance of the RF structure changes, which in turn changes the resonant frequency of the RF structure. In some implementations, during each quarter cycle of the motor, the blades 1310 mesh with the each other. The capacitance of the RF structure increases and the resonant frequency decreases. The process reverses as the blades 1310 unmesh. As a result, the power required to generate the high voltage applied to the dee plate 103 and necessary to accelerate the beam can be reduced by a large factor. In some implementations, the shape of the blades 1310 is machined to form the required dependence of resonant frequency on time.

[00271] The RF frequency generation is synchronized with the blade rotation by sensing the phase of the RF voltage in the resonator, keeping the alternating voltage on the dee plates close to the resonant frequency of the RF cavity. (The dummy dee is grounded and is not shown in Fig. 28).

[00272] The variable reactive element 1308 can be a capacitor formed by a plate 1312 and a surface 1316 of the inner conductor 1300. The plate 1312 is

movable along a direction 1314 towards or away from the surface 1316. The capacitance of the capacitor changes as the distance D between the plate 1312 and the surface 1316 changes. For each frequency range to be swept for one particle energy, the distance D is at a set value, and to change the frequency range, the plate 1312 is moved corresponding to the change in the energy of the output beam.

[00273] In some implementations, the inner and outer conductors 1300, 1302 are formed of a metallic material, such as copper, aluminum, or silver. The blades 1310 and the plate 1312 can also be formed of the same or different metallic materials as the conductors 1300, 1302. The coupling device 1304 can be an electrical conductor. The variable reactive elements 1306, 1308 can have other forms and can couple to the dee plate 100 in other ways to perform the RF frequency sweep and the frequency range alteration. In some implementations, a single variable reactive element can be configured to perform the functions of both the variable reactive elements 1306, 1308. In other implementations, more than two variable reactive elements can be used.

[00274] The control of the gantry, the patient support, the active beam shaping elements, and the synchrocyclotron to perform a therapy session is achieved by appropriate therapy control electronics (not shown).

[00275] Control of the particle therapy system described herein and its various features may be implemented using hardware or a combination of hardware and software. For example, a system like the ones described herein may include various controllers and/or processing devices located at various points. A central computer may coordinate operation among the various controllers or processing devices. The central computer, controllers, and processing devices may execute various software routines to effect control and coordination of testing and calibration.

[00276] System operation can be controlled, at least in part, using one or more computer program products, e.g., one or more computer program tangibly embodied in one or more non-transitory machine-readable media, for execution by, or to control the operation of, one or more data processing apparatus, e.g., a programmable processor, a computer, multiple computers, and/or programmable logic components.

[00277] A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a network.

[00278] Actions associated with implementing all or part of the operations of the particle therapy system described herein can be performed by one or more programmable processors executing one or more computer programs to perform the functions described herein. All or part of the operations can be implemented using special purpose logic circuitry, e.g., an FPGA (field programmable gate array) and/or an ASIC (application-specific integrated circuit).

[00279] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read-only storage area or a random access storage area or both. Elements of a computer (including a server) include one or more processors for executing instructions and one or more storage area devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from, or transfer data to, or both, one or more machine-readable storage media, such as mass PCBs for storing data, e.g., magnetic, magneto-optical disks, or optical disks. Non-transitory machine-readable storage media suitable for embodying computer program instructions and data include all forms of non-volatile storage area, including by way of example, semiconductor storage area devices, e.g., EPROM, EEPROM, and flash storage area devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks.

[00280] Any “electrical connection” as used herein may imply a direct physical connection or a wired or wireless connection that includes intervening components but that nevertheless allows electrical signals to flow between connected components. Any “connection” involving electrical circuitry mentioned herein, unless

stated otherwise, may be an electrical connection and not necessarily a direct physical connection regardless of whether the word “electrical” is used to modify “connection”.

[00281] Any two more of the foregoing implementations may be used in an appropriate combination in an appropriate particle accelerator (e.g., a synchrocyclotron). Likewise, individual features of any two more of the foregoing implementations may be used in an appropriate combination.

[00282] Elements of different implementations described herein may be combined to form other implementations not specifically set forth above. Elements may be left out of the processes, systems, apparatus, etc., described herein without adversely affecting their operation. Various separate elements may be combined into one or more individual elements to perform the functions described herein.

[00283] The example implementations described herein are not limited to use with a particle therapy system or to use with the example particle therapy systems described herein. Rather, the example implementations can be used in any appropriate system that directs accelerated particles to an output.

[00284] Other implementations not specifically described herein are also within the scope of the following claims.

[00285] What is claimed is:

1. A method comprising:
receiving, from a treatment planning process, information that is based on a dose distribution for an irradiation target; and
configuring an adaptive aperture based on the information.
2. The method of claim 1, wherein the information identifies structures of the adaptive aperture; and
wherein configuring the adaptive aperture comprises moving the structures of the adaptive aperture to produce a trimming curve.
3. The method of claim 1, wherein the information identifies characteristics of a trimming curve; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce the trimming curve.
4. The method of claim 1, wherein the information identifies characteristics of pre-trimmed spots, the characteristics being stored in a library in computer memory; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to trim spots of a particle beam so that the spots of the particle beam approximate the pre-trimmed spots.
5. The method of claim 1, wherein the information indicates where spots are to be formed on the irradiation target; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture based on where the spots are to be formed on the irradiation target.
6. The method of claim 1, wherein the information identifies characteristics of spots; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture on a spot-by-spot basis.

7. The method of claim 6, wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture so that at least two successive spots have at least one of a different size or a different shape.

8. The method of claim 1, wherein the information comprises specifications of a trimming curve for a radiation field associated with the irradiation target; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce the trimming curve.

9. The method of claim 8, wherein the specifications identify at least one of shape or a location of the trimming curve relative to a known point.

10. The method of claim 1, wherein the information comprises specifications of distinct trimming curves for different layers of the irradiation target; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce a trimming curve for each layer of the irradiation target based on specifications for the each layer.

11. The method of claim 1, wherein the information comprises specifications of a single trimming curve for a radiation field associated with the irradiation target; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce the single trimming curve, the single trimming curve being used to trim a radiation field for multiple layers of the irradiation target.

12. The method of claim 1, wherein the information comprises specifications for a configuration of the adaptive aperture for a radiation field of the irradiation target; and
wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture based on the specifications for the configuration of the adaptive aperture.

13. The method of claim 12, wherein the information comprises specifications of configurations of the adaptive aperture for different layers of the irradiation target; and wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture for the different layers of the irradiation target based on the information.

14. The method of claim 1, wherein the adaptive aperture comprises:
structures that are movable relative to the irradiation target to produce a shape to trim part of a treatment area; and
a primary carriage coupled to the structures to move the structures relative to the irradiation target in a first dimension relative to the irradiation target; and
wherein configuring the adaptive aperture comprises moving at least one of a structure among the structures or the primary carriage.

15. The method of claim 14, wherein the adaptive aperture comprises secondary carriages coupled to the primary carriage and to the structures to move the structures in a second dimension relative to the irradiation target; and
wherein configuring the adaptive aperture comprises moving at least one of the secondary carriages.

16. The method of claim 15, wherein the second dimension comprises movement into, and out of, the treatment area.

17. The method of claim 15, wherein moving at least one of the secondary carriages comprises rotating the at least one secondary carriage.

18. The method of claim 14, wherein the structures comprise leaves that are extendible and retractable relative to the particle beam.

19. The method of claim 1, wherein configuring the adaptive aperture comprises rotating all or part of the adaptive aperture.

20. A method comprising:
receiving, from a treatment planning process, information that is based on a dose distribution for an irradiation target; and
performing at least one of the following operations:
moving structures to trim spots of a particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received;
moving structures to produce a trimming curve for a layer of an irradiation target based on a specification of a trimming curves for the layer included in the information received;
moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included in the information received; or
moving structures based on configuration information for the structures in the information received.

21. The method of claim 20, wherein the structures are part of an adaptive aperture and comprise leaves that are movable into, or out of, an area to be treated by the particle beam.

22. A particle therapy system comprising:
a particle accelerator to output a particle beam for application to an irradiation target;
an adaptive aperture between the particle accelerator and the irradiation target, the adaptive aperture comprising structures that are movable into, and out of, a path of the particle beam; and
a control system comprising one or more processing devices configured to perform operations comprising:

receiving, from a treatment planning process, information that is based on a dose distribution for the irradiation target; and
configuring the adaptive aperture based on the information.

23. The particle therapy system of claim 22, wherein the information identifies characteristics of pre-trimmed spots, the characteristics being stored in a library in computer memory; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to trim spots of a particle beam so that the spots of the particle beam approximate the pre-trimmed spots.

24. The particle therapy system of claim 23, wherein the information indicates where the pre-trimmed spots are to be formed on the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture based on where the pre-trimmed spots are to be formed on the irradiation target.

25. The particle therapy system of claim 22, wherein the characteristics comprise one or more parameters, the one or more parameters comprising at least one of a radius associated with a spot or an angle of rotation of the spot.

26. The particle therapy system of claim 22, wherein the information identifies characteristics of spots; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture on a spot-by-spot basis.

27. The particle therapy system of claim 26, wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture so that at least two successive spots have at least one of a different size or a different shape.

28. The particle therapy system of claim 26, wherein configuring the adaptive aperture comprises maintaining structures of the adaptive aperture so that at least two successive spots have at least one of a same size or a same shape.

29. The particle therapy system of claim 22, wherein the information comprises specifications of a trimming curve for a radiation field associated with the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce the trimming curve.

30. The particle therapy system of claim 29, wherein the specifications identify at least one of shape or a location of the trimming curve relative to a known point.

31. The particle therapy system of claim 22, wherein the information comprises specifications of trimming curves for different layers of the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce a trimming curve for each layer of the irradiation target based on specifications for the each layer.

32. The particle therapy system of claim 22, wherein the information comprises specifications of a single trimming curve for a radiation field associated with the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture to produce the single trimming curve, the single trimming curve being used to trim a radiation field for multiple layers of the irradiation target.

33. The particle therapy system of claim 22, wherein the information comprises specifications for a configuration of the adaptive aperture for a radiation field of the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture based on the specifications for the configuration of the adaptive aperture.

34. The particle therapy system of claim 33, wherein the information comprises specifications of configurations of the adaptive aperture for different layers of the irradiation target; and

wherein configuring the adaptive aperture comprises moving structures of the adaptive aperture for the different layers of the irradiation target based on the information.

35. The particle therapy system of claim 22, wherein the adaptive aperture comprises a primary carriage coupled to the structures to move the structures relative to the irradiation target in a first dimension relative to the irradiation target; and

wherein configuring an adaptive aperture comprises moving at least one of a structure among the structures or the primary carriage.

36. The particle therapy system of claim 35, wherein the adaptive aperture comprises secondary carriages coupled to the primary carriage and to the structures to move the structures in a second dimension relative to the irradiation target; and

wherein configuring an adaptive aperture comprises moving at least one of the secondary carriages.

37. The particle therapy system of claim 36, wherein the second dimension comprises movement into, and out of, the treatment area.

38. The particle therapy system of claim 36, wherein moving at least one of the secondary carriages comprises rotating the at least one secondary carriage.

39. The particle therapy system of claim 35, wherein the structures comprise leaves that are extendible and retractable relative to the particle beam.

40. The particle therapy system of claim 22, wherein configuring the adaptive aperture comprises rotating all or part of the adaptive aperture.

41. A particle therapy system comprising:
a particle accelerator to output a particle beam for application to an irradiation target; and
a control system comprising one or more processing devices configured to perform operations comprising:
 receiving, from a treatment planning process, information that is based on a dose distribution for the irradiation target; and
 performing at least one of the following operations:
 moving structures to trim spots of the particle beam so that the spots of the particle beam approximate pre-trimmed spots for which characteristics are obtained based on the information received;
 moving structures to produce a trimming curve for a layer of an irradiation target based on a specification of a trimming curve for the layer included in the information received;
 moving structures to produce a single trimming curve for all radiation fields of an irradiation target based on specifications of the single trimming curve included in the information received; or
 moving structures based on configuration information for the structures in the information received.

42. The particle therapy system of claim 41, wherein the structures are part of an adaptive aperture and comprise leaves that are movable into, or out of, an area to be treated by the particle beam.

43. A method comprising:
obtaining a target dose distribution for a treatment area;

identifying one or more pre-trimmed spots from among a library of pre-trimmed spots that best approximates one or more points of the target dose distribution; and outputting information identifying the one or more pre-trimmed spots.

44. The method of claim 43, further comprising:
obtaining a set of coarse spots that approximates the target dose distribution;
changing an intensity of one or more spots in the coarse set based on a maximum acceptable dosage for a portion of the treatment area; and
identifying a spot among the coarse spots that is within a transition region of the treatment area;
wherein identifying comprises substituting the one or more pre-trimmed spots for the spot identified.

45. The method of claim 44, further comprising:
obtaining a minimum acceptable dosage for a portion of the treatment area;
wherein the set of coarse spots is selected based on the minimum acceptable dosage.

46. The method of claim 43, wherein the information is usable to configure an adaptive aperture to achieve a dose distribution during treatment that approximates the target dose distribution.

47. One or more non-transitory machine-readable storage media storing instructions that are executable to perform operations comprising:
obtaining a target dose distribution for a treatment area;
identifying one or more pre-trimmed spots from among a library of pre-trimmed spots that best approximates one or more points of the target dose distribution; and
outputting information identifying the one or more pre-trimmed spots.

48. The one or more non-transitory machine-readable storage media of claim 47, wherein the operations comprise:

obtaining a set of coarse spots that approximates the target dose distribution;
changing an intensity of one or more spots in the coarse set based on a maximum acceptable dosage for a portion of the treatment area;
identifying a spot among the coarse spots that is within a transition region of the treatment area; and
substituting the one or more pre-trimmed spots for the spot within the transition region.

49. The one or more non-transitory machine-readable storage media of claim 48, wherein the operations comprise:

obtaining a minimum acceptable dosage for a portion of the treatment area;
wherein the set of coarse spots is selected based on the minimum acceptable dosage.


50. The one or more non-transitory machine-readable storage media of claim 47, wherein the information is usable to configure an adaptive aperture to achieve a dose distribution during treatment that approximates the target dose distribution.

51. A method comprising:
obtaining treatment planning information;
determining that an irradiation target has moved since a prior application of radiation based on the treatment planning information;
modifying the treatment planning information based on movement of the irradiation target to produce updated treatment planning information; and
applying radiation to the irradiation target based on the updated treatment planning information.

52. The method of claim 51, wherein modifying the treatment planning information is performed by a computing system of a particle therapy system.

53. The method of claim 51, wherein determining that that irradiation target has moved is based on a comparison of images of the irradiation target captured between two points in time.

54. The method of claim 51, wherein the treatment planning information is modified by changing locations to which spots of radiation are to be applied to the irradiation target based on a location of the irradiation target following movement.



2/37




FIG. 2




FIG. 4




FIG. 3

4/37




FIG. 5




FIG. 6

6/37




FIG. 7




FIG. 8

8/37




FIG. 9

9/37




FIG. 10

10/37




FIG. 11

11/37




FIG. 12




FIG. 13

12/37




FIG. 14




FIG. 15

14/37




FIG. 16

15/37




FIG. 17




FIG. 18

17/37




FIG. 19

18/37




FIG. 20

19/37




FIG. 21

20/37




FIG. 22

21/37




FIG. 24




FIG. 23

22/37




FIG. 25

23/37




FIG. 26




FIG. 27




FIG. 28




FIG. 29




FIG. 30

28/37




FIG. 31

29/37




FIG. 32

30/37




FIG. 33




FIG. 34

32/37




FIG. 35

33/37




FIG. 36

34/37




FIG. 37

35/37




FIG. 38




FIG. 39




FIG. 40




FIG. 41




FIG. 42

37/37




FIG. 43

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/041076

A. CLASSIFICATION OF SUBJECT MATTER
INV. A61N5/10
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
A61N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2015/003111 A1 (UNIV IOWA RES FOUND [US]) 8 January 2015 (2015-01-08) paragraphs [0002], [0078], [0097], [0098], [0102], [0120] - [0133]; claims 1,13,14,17; figures 2,3,10-13 -----	22-42, 47-50



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 September 2017

Date of mailing of the international search report

27/09/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Kajzar, Anna

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/041076

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☒ Claims Nos.: 1-21, 43-46, 51-54
because they relate to subject matter not required to be searched by this Authority, namely:
Rule 39.1(iv) PCT - Method for treatment of the human or animal body by therapy and surgery
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/041076

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2015003111 A1	08-01-2015	CN 105359223 A	24-02-2016
		EP 3014629 A1	04-05-2016
		JP 2016526452 A	05-09-2016
		US 2016199667 A1	14-07-2016
		WO 2015003111 A1	08-01-2015
