Superconducting Modular 10 MW AI PODs

Peter Abrell
Vision Electric Superconductors
abrell@vesc-superbar.de

Table Of Contents

1. **Introduction** 1
  1.1. **Background** 1
  1.2. **Specification Scope and Goals** 1
2. **Data Center Electrical Infrastructure** 3
  2.1. **MV Switchyard** 3
  2.2. **Onsite AC-Generation** 4
  2.3. **AC/DC Conversion** 4
  2.4. **BESS and DC onsite generation** 5
  2.5. **HTS System** 6
    2.5.1. **Busbar** 7
    2.5.2. **Main Current Leads** 10
    2.5.3. **Tap-Off Boxes** 11
    2.5.3.1. **Tap-off Sidecar (TOS)** 12
    2.5.3.2. **Tap-off to Rack (TOR)** 14
    2.5.4. **Cooling Systems** 16
  2.6. **GPU POD** 17
  2.7. **Auxiliary Power** 18
3. **Benefits and Concluding Remarks** 19

Acronyms

AC - Alternating current

AFE - Active front end

BESS - Battery energy storage system

CDU - Cooling distribution unit

DC - Direct current

FC - Fuel cell

HTS - High-temperature superconductors

HV - High voltage

MV - Medium voltage

POD - Performance optimized data center

PV - Photovoltaic

UPS - Uninterruptable power supply

1. Introduction

1.1. Background

The next generation of AI data centers will consist of racks with power ratings of 1 MW and more. Accommodating this computing load will require a new type of electrical and thermal systems. The following specification describes an AI factory the utilizes sustainable building materials, liquid cooling and high-temperature superconductors (HTS). It represents a new approach to building data centers that focuses on the unique benefits of superconductors.

When cooled below their critical temperature, HTS exhibit no direct current (DC) resistance. Ohmic losses are traded for a base cooling load at cryogenic temperatures (77 K, -196°C). The current carrying capacity of an HTS-conductor is defined by:

- Operating temperature (65 K to 77 K with liquid nitrogen)
- Surrounding magnetic field

In a conventional system, to transfer more power the voltage level is increased to limit the ohmic losses and conductor cross-section. Since HTS do not have ohmic losses, higher voltage levels are not needed. The distribution voltage can be set at 800 VDC. For MW-sized performance optimized datacenters (POD) this results in operating currents in the range of kilo amperes. For superconductors these currents are feasible. Due to the small conductor size and low voltage, the distance between the conductors is small, reducing the resulting magnetic field and resulting forces.

The specification is set up in a modular way because:

- HTS and the 800 VDC distribution voltage are not yet used in data centers, this modular design can be used as proof of concept for a single POD project combining both technologies.
- Cooling and power systems are designed according to best practices
- AC/DC rectifiers can supply this power in one transformer-rectifier unit.
- Redundancy can be implemented on a per POD basis with two distribution systems, A and B, allowing high reliability.
- DC-based components like batteries and fuel cells can be implemented on the DC side and allow for higher efficiency

1.2. Specification Scope and Goals

The overall scope of the specification starts at the high voltage (HV) connection to the medium voltage switchyard (MV). The following technologies will be

- AC/DC converters
- battery storage systems (BESS)
- on-site generation, AC and DC based options
- cryogenic cooling system
- current lead to HTS system
- HTS busbar
- tap-off boxes

The specification defines the:

- layout of a data center consisting of multiple PODs

- electrical and mechanical infrastructure
- cooling systems
    - for IT-hardware
    - for the HTS busbar
- redundancy concept

It will showcase the:

- implementation of HTS technology
- implementation of 800 VDC, from rectification to rack
- redundancy concept

2. Data Center Electrical Infrastructure

Power is supplied to the data center via an HV grid connection and switchyard. The incoming voltage is stepped down to MV to supply rectifier units, BESS and on-site generation. The rectifiers convert AC power into DC for downstream distribution within the data center.

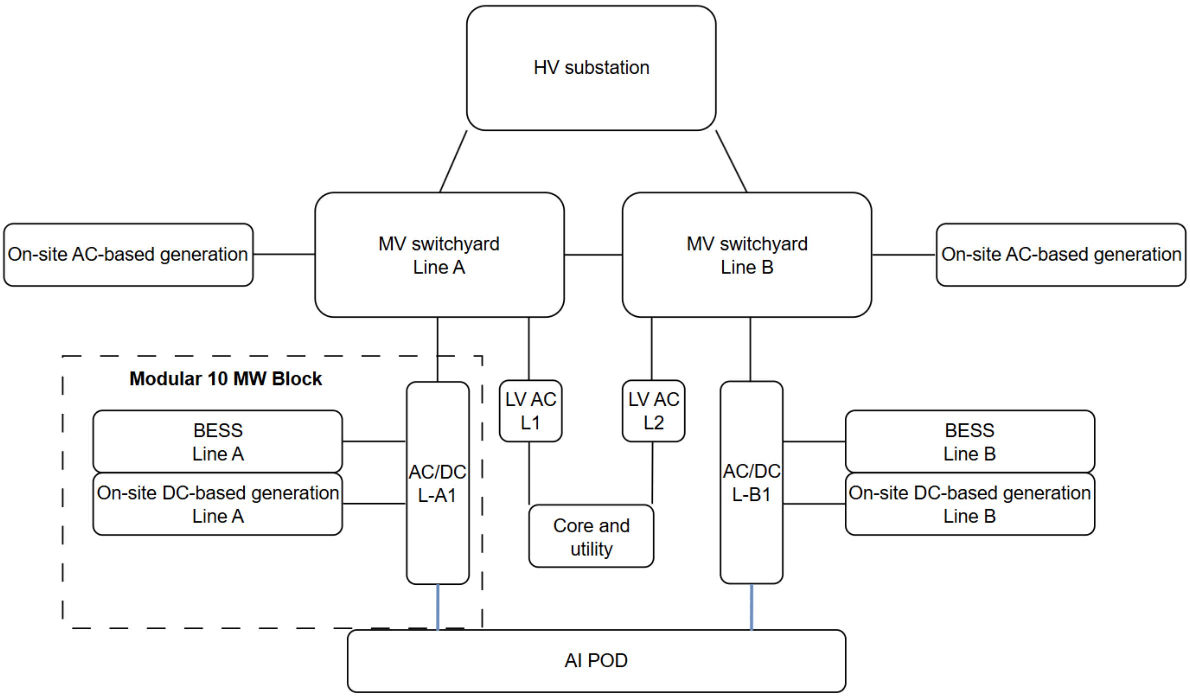

Figure 1 Overview of the main electrical components

2.1. MV Switchyard

The MV switchyard is implemented in a fully redundant configuration to supply both Line A and Line B with uninterrupted power in the event of a switchyard failure. Each MV switchyard is rated for twice the total power demand of all POD modules in the data center, ensuring full load coverage under single-fault conditions.

The MV switchyards are interconnected via circuit breakers to protect against failures at the HV switchyard level. This arrangement provides redundancy at both the MV and HV levels. Utility power is used to supply data center auxiliary loads and core infrastructure functions.

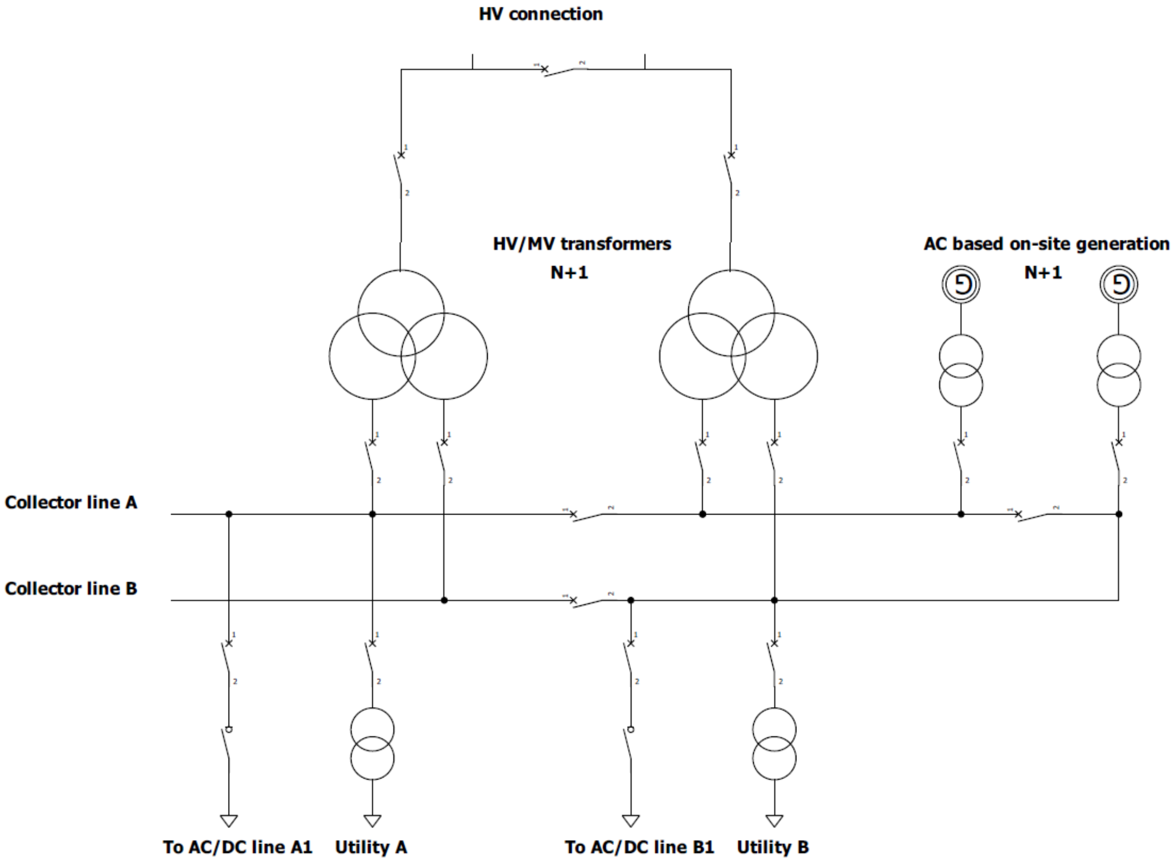

Figure 2 Basic single line diagram of the medium voltage switchyard.

2.2. Onsite AC-Generation

On-site AC generation based on rotating machines is connected to the MV switchyard and implemented with per-line redundancy at an N+1 level. The choice of generator technology depends on local availability and the primary energy source. Gas generators operating on natural gas can be deployed and designed to be future proof by enabling operation with hydrogen. Diesel generators represent an alternative, particularly where gas infrastructure is unavailable.

As an alternative to rotating machines, large-scale battery energy storage systems can be used to replace conventional generators. As these technologies are inherently DC-based, they are connected directly to the DC bus. Optionally, they can be implemented at the MV level to supply the entire data center via centralized power conversion, rather than being distributed across individual modular units.

2.3. AC/DC Conversion

The modular AC/DC conversion units represent one of the key architectural differences compared to conventional data centers. These units convert medium-voltage AC into 800 V DC, which is then supplied to the AI POD via the superconducting distribution line.

Implementing a common DC bus that interconnects multiple distribution lines contradicts the modular design philosophy and introduces additional challenges, including system security, DC circuit breaker coordination, and converter control stability when multiple units are connected without sufficient impedance between them. To avoid a large, shared DC bus, battery storage and on-site DC generation are integrated in a modular manner on a per distribution line basis.

Industrial rectifier systems are well established and widely used in applications requiring very high currents, such as electrolysis. In industrial environments, 12-pulse diode rectifiers are commonly employed. These systems typically operate with a current control and relatively constant loads. While they exhibit a comparatively high DC ripple (~10%), this is generally acceptable for industrial processes.

In contrast, the power electronic infrastructure of AI data centers requires a much cleaner DC to supply the converters downstream. Furthermore, the highly dynamic load profiles of AI data centers demand fast-responding, actively controllable rectifier systems rather than conventional diode-based solutions. Power swings can range from 30% to 100% load, in the space of 100 ms. For a 200 kW rack, this results in worst-case ramps of up to 1.75 A/ms.

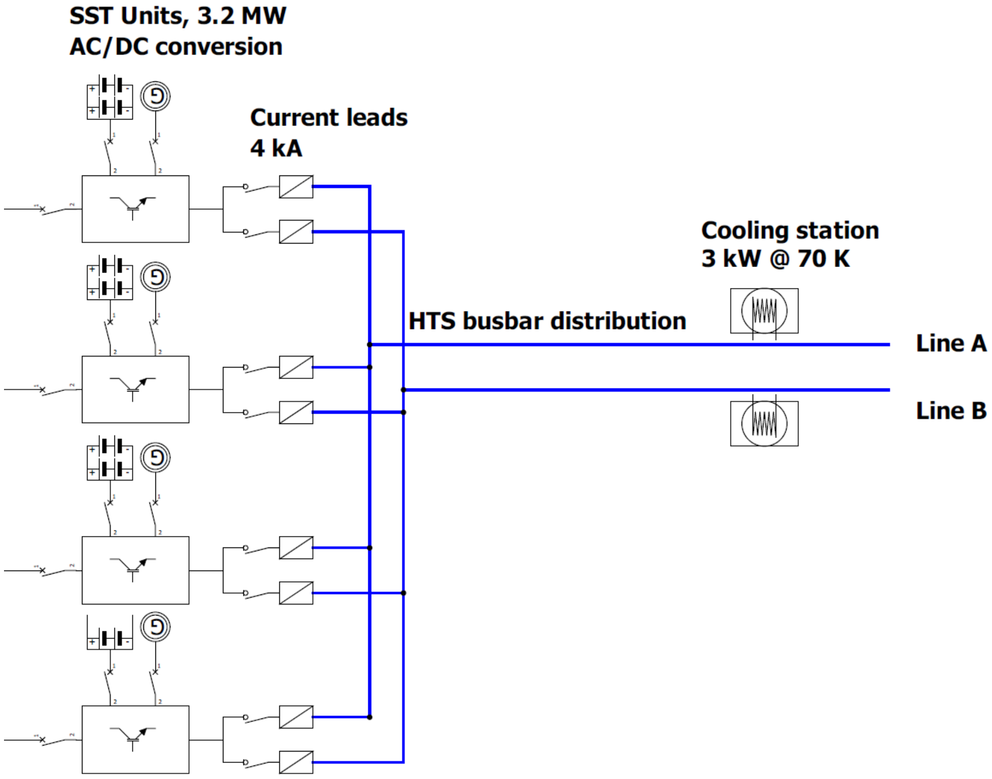

Figure 3 Single line diagram of the 10 MW modular set up. It includes the connection to the MV, to recitification all the way down to the IT load.

Solid state transformers (SSTs) are best suited for this application. They provide significantly lower DC ripples while requiring substantially less reactive power from the grid. The rectifier is implemented as a skid-mounted system that includes an MV/LV transformer. Converters of this type are commercially available and can be delivered with lead times of less than one year. A 4 to make 3 redundancy concept is used.

2.4. BESS and DC onsite generation

Batteries and accumulators are inherently DC-based and are already widely deployed at an 800 V level in the electric vehicle market. They can be used as a flexible energy reserve or simply back-up power. In combination with PV systems, they supplement on-site generation and increase the share of renewable energy. Operationally, the BESS system will be used to buffer load swings.

Battery capacity is selected based on the installed PV power and the required bridging time until full operation of the backup generation system. The battery systems can also be used to relieve the grid and allow the operator to provide flexibility. Battery energy storage systems in the multi-MW and MWh range have already been deployed by several vendors, for example the Tesla Megapack, or container solutions by Tesvolt. They are placed as close as possible to the infeed of the superconducting system to reduce the length of the DC bus.

Connecting elements and terminations for a battery-DC bus in the MW-range already exist, see for example the megawatt charging system: a low voltage, water-cooled (<1.5 kV) DC cable system. Alternatively, copper busbars can be used. By parallelizing battery units, the size of the protection gear can be reduced.

The 800 V bus allows the concept to be future proof for green hydrogen economy. Instead of conventional gas or diesel generators, fuel cells can be incorporated as backup power sources, further reducing carbon emissions. A challenge lies in ensuring a reliable hydrogen supply infrastructure. NorthC installed a 500 kW fuel cell system to supplement the back-up power supply in a data center in Groningen, Netherlands. However, systems in the multi-MW-range still need to be implemented. [https://www.northcdatacenters.com/en/news/northc-installs-europes-first-emergency-power-facilities-that-run-on-green-hydrogen-in-groningen/].

2.5. HTS System

The HTS system consists of the busbar, carrying the HTS conductor and liquid nitrogen, the current lead and tap-off boxes that act as an infeed and out feed, and the cooling machines that keep the circulating liquid nitrogen at operating temperature.

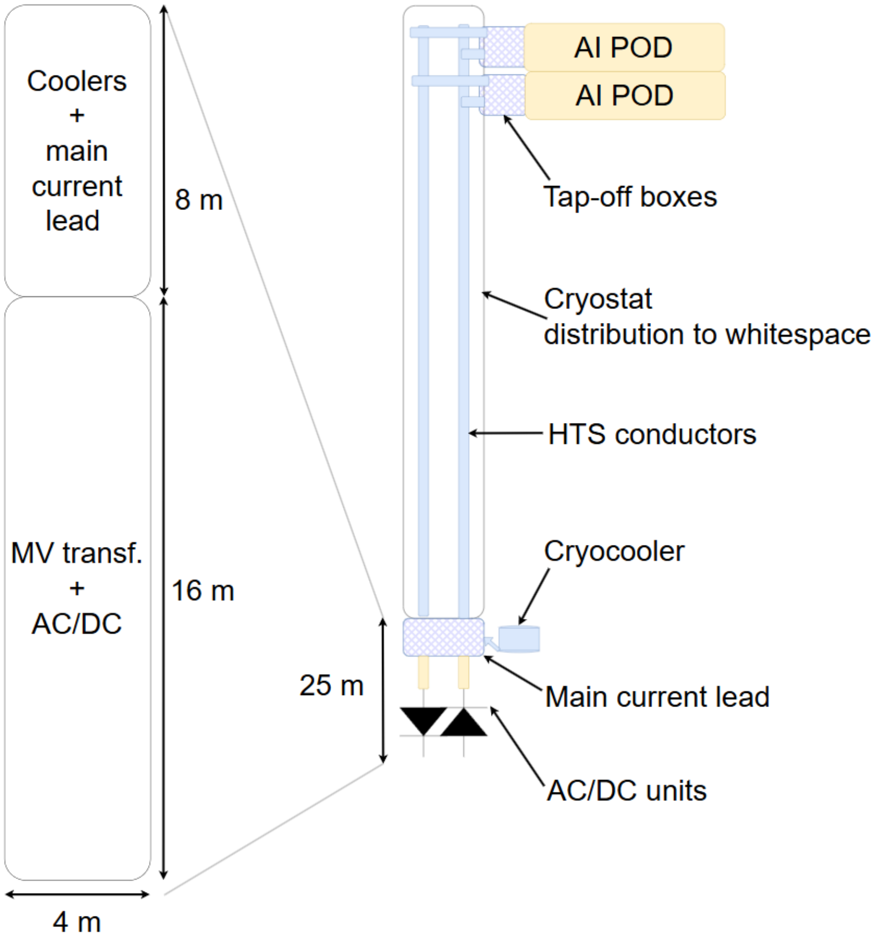

Figure 4 Overview of the basic elements making up the HTS system. Multiple conductors can be placed inside a cryostat.

Two lines are used per module, allowing for redundancy in the distribution system. The cooling machines used to circulate and cool the liquid nitrogen are set-up with an n+2 redundancy to allow for redundancy while the machines are being maintained.

HTS distribution lines can be routed either above or below the racks. In the event of a vacuum leak in a cryostat, water may condense on the affected vacuum pipe. To mitigate this risk, cryostats can be installed in an underfloor area with a drain. The HTS busbar has low space requirements and generate no heat, so it is possible to accommodate MW distribution in compact areas such as underfloor channels. The cryostats fully encapsulate the conductors, and due to the low voltage and minimal spacing between them, electromagnetic emissions are negligible.

The system consists of the HTS busbar, current lead, tap-offs and cooling management. They will be described in the next chapters.

2.5.1. Busbar

The HTS busbar combines the needed components to make the HTS system work. It houses the conductors which are submerged in liquid nitrogen. To make the system efficient and limit the consumption of liquid nitrogen, it must also be thermally insulated. The next sections detail the essential elements.

2.5.1.1. HTS Conductor

The conductor consists of HTS tapes, copper and solder. HTS by themselves are mechanically fragile and difficult to handle. By incorporating them into a composite conductor the drawbacks of the tapes are minimized. The current carrying capacity of the tapes is dependent on the operating temperature and the magnetic field that the tape is affected by, Figure 5 shows this relationship.

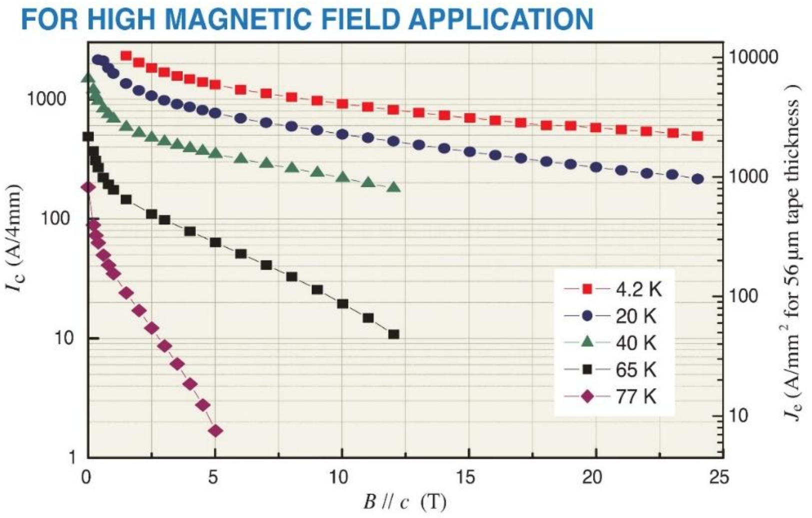

Figure 5 Critical current of a 4 mm wide HTS tape [FFJ].

Assuming the AI POD consists of some of the newest Nvidia GPUs, multiple Vera Rubin NVL576 racks (rack power of up to 600 kW) require 12.8 MW per POD. This results in a maximum operational current of 16 kA. The current is divided by 4 AC/DC units, resulting in a conductor with an operational current of 4 kA.

Table 1 ICE® Wire specifications and dimensions. Right: Image of an HTS composite conductor.

Max. Operational Current (77 K, s.f.)	4.0 kA
Critical Current (77 K, s.f.)	5.8 kA
Number of 12 mm HTS tapes	11
Thickness HTS tape	0.1 mm
Number of copper tapes	12
Thickness copper tape	0.1 mm
Conductor thickness	3 mm

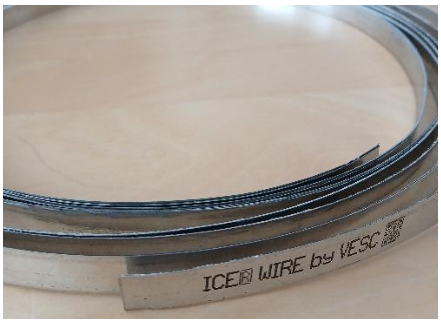

The fault behavior of HTS systems is unique in electrical engineering. Due to their very low impedance, they can initially generate high fault currents. However, HTS conductors inherently exhibit fault current limiting characteristics. When the current nears the critical current of the HTS tapes, an increasing share of the current is transferred to the surrounding copper and solder, which increases the overall electrical resistance of the conductor and provides a current limiting effect.

This transition leads to increased heat generation, raising the conductor temperature and further increasing the current through copper and solder. As a result, system resistance continues to increase, further reducing the fault current. It is essential to clear the fault before the temperature rises to a level that could damage the HTS tapes. To extend the allowable fault clearance time, additional copper can be incorporated into the HTS composite conductor.

2.5.1.2. Cryostat

The cryostat acts as the thermal insulation and prevents the liquid nitrogen from heat input. It is needed to ensure the efficient operation of the busbar. Incase multiple modules are used; there are also multiple conductors. In this case they can share a common cryostat. One distribution line would then share a cryostat which would branch out to the individual PODs. A cryostat with an outer diameter of 260 mm is sufficient to easily carry 5 PODs or distribute 50 MW. Linkages are possible by using T-segments.

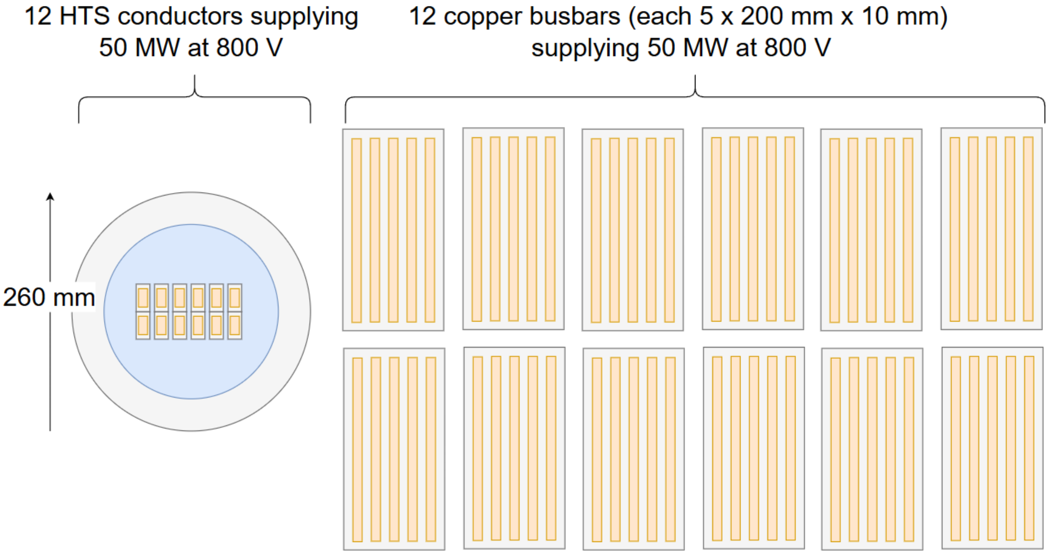

Figure 6 Comparison of an HTS and a copper system supplying 50 MW at 800 V.

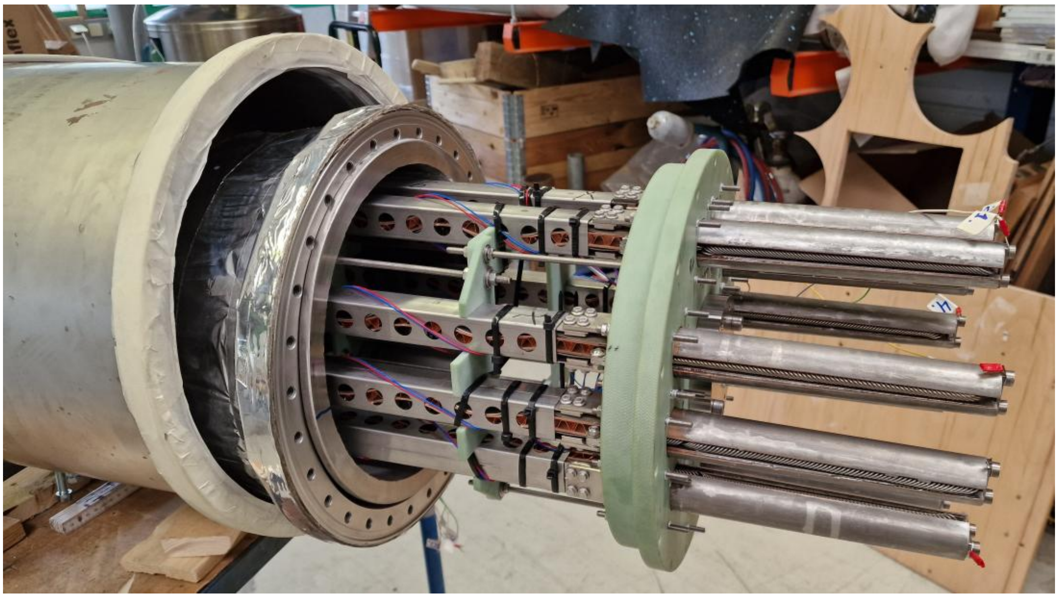

Figure 7 Busbar used in the Demo200 project. 10 individual busbars each rated at 20 kA in a single cryostat. This would equal 160 MW at 800 V.

Cryostat costs can be significantly reduced through scaling. Typical costs are approximately 1,000 € per meter. For a 10 MW system, this corresponds to about 100 € per MWm. If capacity is increasd to 50 MW, costs decrease to roughly 20 € per MWm, as multiple lines can be integrated into a single cryostat.

2.5.2. Main Current Leads

The main current leads represent the infeed to the HTS system. Figure 8 shows a gas-cooled current lead with a rated current of 20 kA.

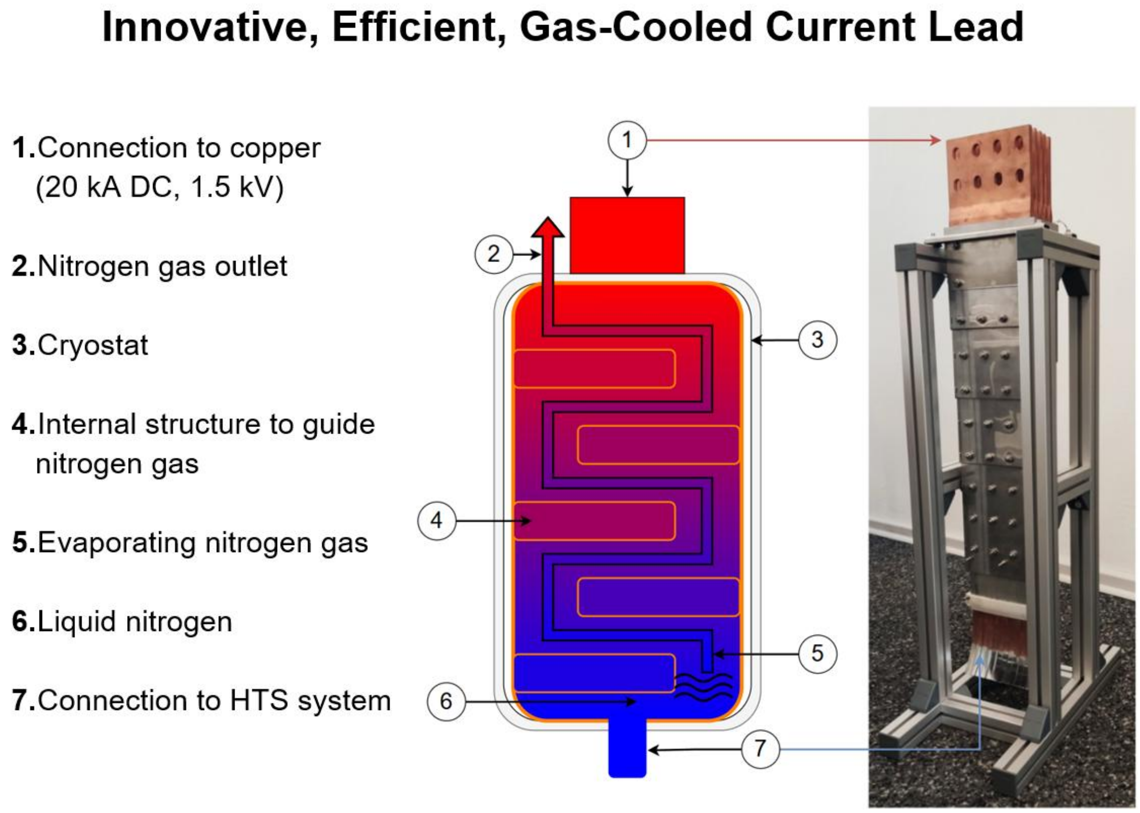

Figure 8 Left: overview of a gas-cooled current lead for 20 kA. Right: tap-off box for the outfeed.

Figure 9 20 kA current lead used in the 3S project. Left: current lead with copper connections at the top, HTS connection at the bottom. Right: compressor used for cryocooler.

The ratio between cross section and length of current lead (copper) needs to be at a certain ratio to ensure efficiency. There are two main forms of heat entry that need to be cooled:

- Heat conduction through the copper
- Heat generated by I^2R losses of the copper

Increasing the length of the current lead reduces heat conduction but increases I^2R losses. The same is inversely true for the cross section: a bigger cross section reduces the resistance and so electric losses but also increases the heat conduction. An optimal geometric ratio exists at which the combined heat load from conduction and electrical losses is minimized. The current defines the cross section of the copper (typically around 1 A/mm^2) and therefore also the length. If needed, current density in the copper can be increased with water cooling (up to around 10 A/mm^2), therefore also decreasing the size of the current lead.

The main current leads will directly connect to the SSTs and will consist of two phases each with a maximum operational current of 4 kA, and cooling requirement of 400 W.

2.5.3. Tap-Off Boxes

The tap-off boxes supplying the IT loads perform the same function as current leads. They form the electrical outfeed of the HTS system and provide the interface to the POD. Their size depends on the selected termination concept to the POD, which can be freely chosen. Two termination variants are discussed in the following.

Termination can be implemented either at the POD level, where the full current is routed through a single tap-off, or at the rack level, where each rack is equipped with its own smaller current lead carrying a reduced current. An in-between solution is also possible, for example feeding two racks with a tap-off. Rack-level termination enables the use of smaller DC circuit breakers but increases the complexity of liquid routing and piping, while larger sidecar tap-offs fault current will be switched by the medium voltage AC breaker before the rectifier.

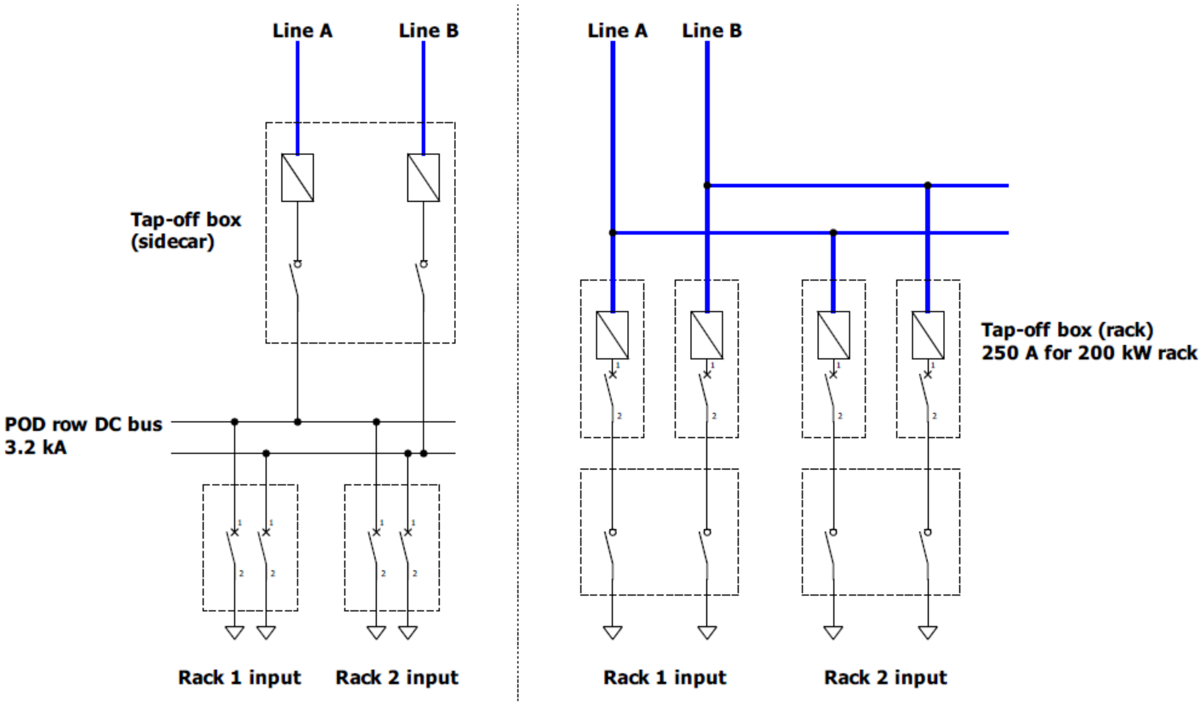

Figure 10 Two possible variations for the tap-off (outfeed of the HTS system) to the IT loads. Left: all of the current is output to a common DC bus that distributes the current in the POD room. Right: the outfeed is distributed to each rack.

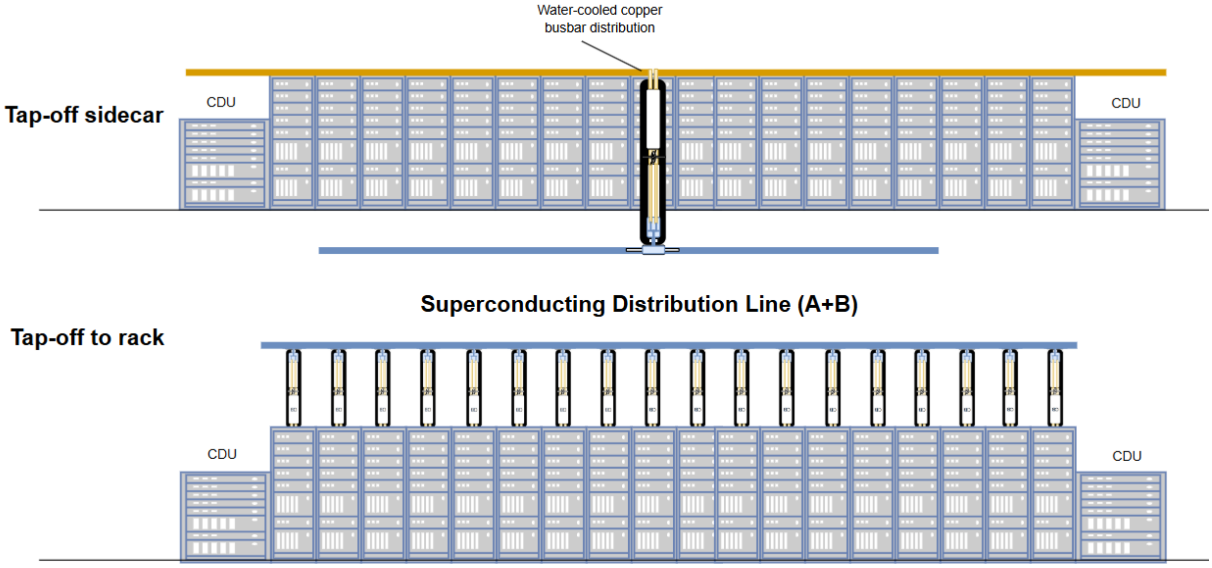

Figure 11 Overview of the two discussed tap-off variants.

2.5.3.1. Tap-off Sidecar (TOS/RPP)

This approach replaces distributed tap-off boxes with a dedicated sidecar that houses the current leads for both the A and B supply lines. It acts as a remote power panel that has switching capabilities. Supply lines A and B can be routed above the racks or inside the data center floor. The current is carried to the racks with a water-cooled busbar to reduce the space requirements in the white room. Fault currents are switched at the rectifier's AC input.

The TOS is the size of a standard IT rack and includes both current lead for A and B lines.

- the current lead (2 kA or 1,600 kW)
    - conduction cooled by liquid nitrogen (50 W/kA)
    - length: 1.0 m
    - cross section: 570 mm^2
    - total cooling power for current lead unit 200 W at 77 K
    - water-cooled copper, current density of 3.5 A/mm^2
- temperature control
    - used for startup and shutdown to limit cold entry intro monitoring and switching systems
- monitoring, switching and CDU systems.
    - monitors cooling systems, water-cooled switch and distribution

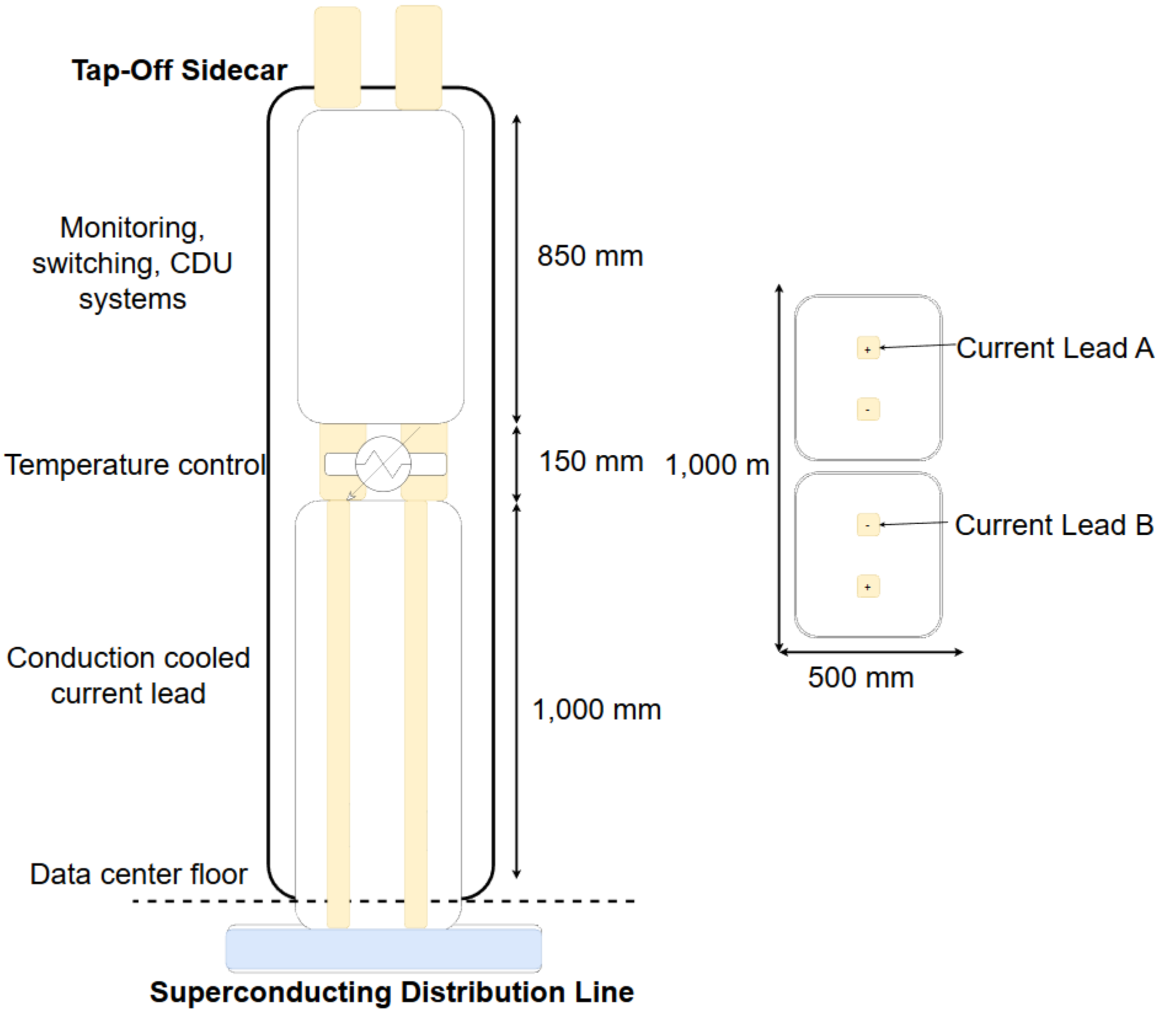

Figure 12 Tap-off sidecar system, acting as a RPP. Left: frontal view. Right: top-down view of the current leads and A, B systems.

DC breakers for this current rating are available from multiple manufacturers, an example can be found in Figure 13.

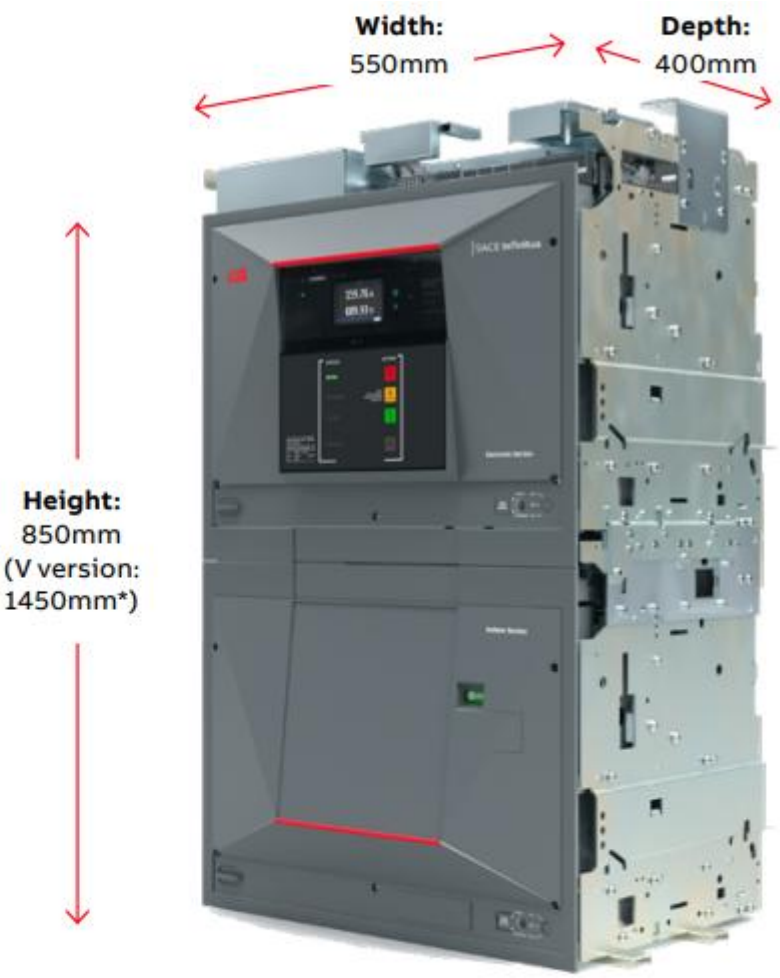

Figure 13 ABB SACE Infinitus breaker for low voltage with a nominal current of 2,500 A.

Sidecar termination could also be designed to comply with the Mt. Diablo 400 specification. It defines an AC/DC sidecar delivering a ±400 VDC output. In this configuration, the tap-off boxes fully replace the sidecar units and supply ±400 VDC from the rectification system. The Mt. Diablo sidecar output is specified as a ±400 VDC busbar with a maximum power of 1.1 MW. Consequently, each POD is supplied by ten tap-off boxes, each delivering up to 1.2 kA at ±400 VDC.

2.5.3.2. Tap-off to Rack (TOR)

The smaller, more distributed tap-offs distribute smaller current (250 A) and therefore have a much higher shape factor than the bigger 3.2 kA sidecar tap-offs, 14,000 m^{-1} compared to 280 m^{-1}. Due to the much smaller current, it is more efficient to have a longer current lead to limit heat conduction. This results in a long, thin current lead. To reduce the length of the current lead the copper terminals will be water cooled.

- the current lead (250 A or 200 kW)
    - conduction cooled by liquid nitrogen (50 W/kA)
    - length: 0.5 m
    - cross section: 36 mm^2

    - water-cooled copper to reduce length, current density of 7 A/mm^2
    - total cooling power for tap-off unit 25 W at 77 K
- temperature control
    - used for startup and shutdown to limit cold entry intro monitoring and switching systems
- DC circuit braker
    - monitors cooling systems, water-cooled switch and distribution

A key challenge in DC distribution is safe interruption of fault currents and fault current interrupting capabilities can be implemented in this tap-off. For nominal currents of 300 A and voltage of up to 1 kV they are available as stock items (https://www.ato.com/300a-dc-circuit-breaker).

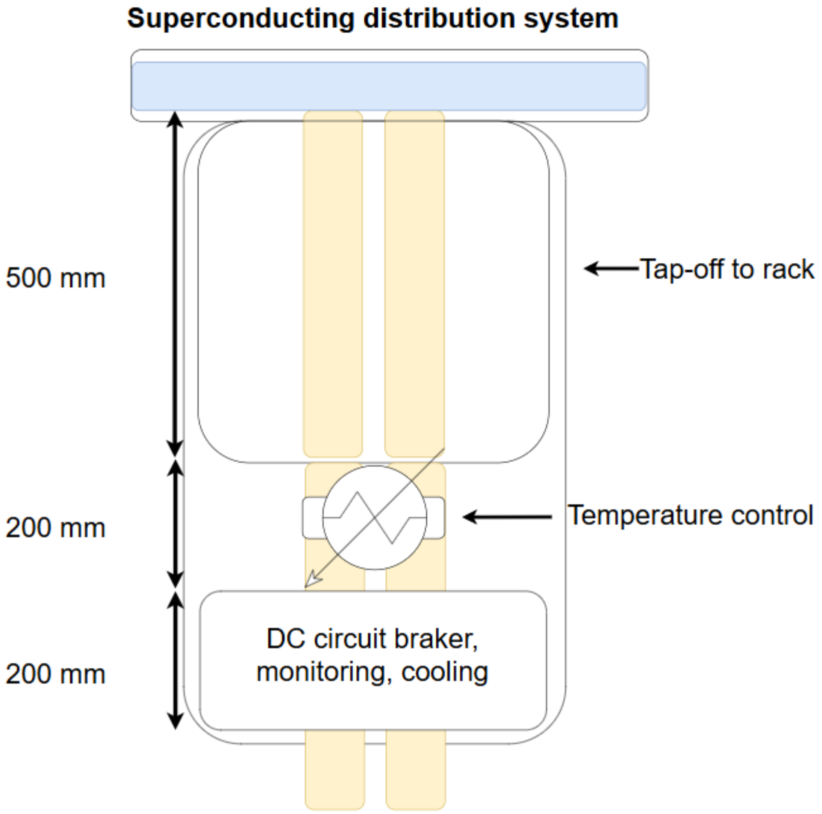

Figure 14 Overview of tap-off to rack system.

300A DC Circuit Breaker, 3 Pole

★★★★★
5 out of 5 based on 1 reviews | Write a review

$255.92

Lower price 300A DC circuit breaker with 20 ms total break time for sale online. Working temperature of smart circuit breaker between -25℃ and 50 ℃. PV DC circuit breaker is commonly used in power transfer switch, overload protection device and control appliance.

Number of Poles * 3 pole, 750V DC

SKU: ATO-DCCB-300A

Figure 15 Website screenshot of 300 A DC breaker that can be used for tap-offs

2.5.4. Cooling Systems

Cooling machines are needed to cool the heat intake of the cryogenic system. The current leads require a specific cooling power of around 50 W/kA at 70 K. The other main heat intake is from the cryostat, around 2 W/m enter through the cryostat. The distribution length totals around 100 m. Two pumps will be used for redundancy, each pump causes roughly 200 W heat entry. Other losses such as hysteresis losses or coupling losses are so low that they can be neglected.

For a 10 MW unit the maximum cooling power results in:

Current lead and tap-offs:	50 W/kA x 16 kA x 4 =	3,200 W
Heat input of cryopumps	2 x 200 W =	400 W
Cryostat heat entry (100m distribution):	2 W/m x 50 m x 2 =	200 W
	Total:	**3,800 W**

Cooling takes place at 70 K and N+2 redundancy is needed in case of breakdown during the maintenance.

These requirements can be full fulfilled by commercially available cooling machines, delivery time for the cold heads and compressors is around 8 – 10 weeks. They require annual maintenance. Several cryocoolers are readily available:

- CryoMech AL-600
    - 500 W at 70 K
    - 9 units to supply 4,500 W at 70 K
    - Installed compressor power of 12.5 kW per cooler, to a total of 112.5 kW at 480 V

- Sumitomo CH-160D3
    - 650 W at 70 K
    - 8 units to supply 5,200 W at 70 K
    - Installed compressor power of 14 kW per cooler, to a total of 98 kW at 480 V

- Stirling SPC-1
    - 1,000 W at 70 K
    - 6 units to supply 6,000 W at 70 K
    - Installed power of 12 kW per cooler, to a total of 72 kW at 480 V for a container solution

To allow for best operational efficiency the Stirling SPC-1 is the best option, while also allowing for highest power during start-up. Including control, pumps and piping the required space is enough to house two cooling stations inside the power conversion room.

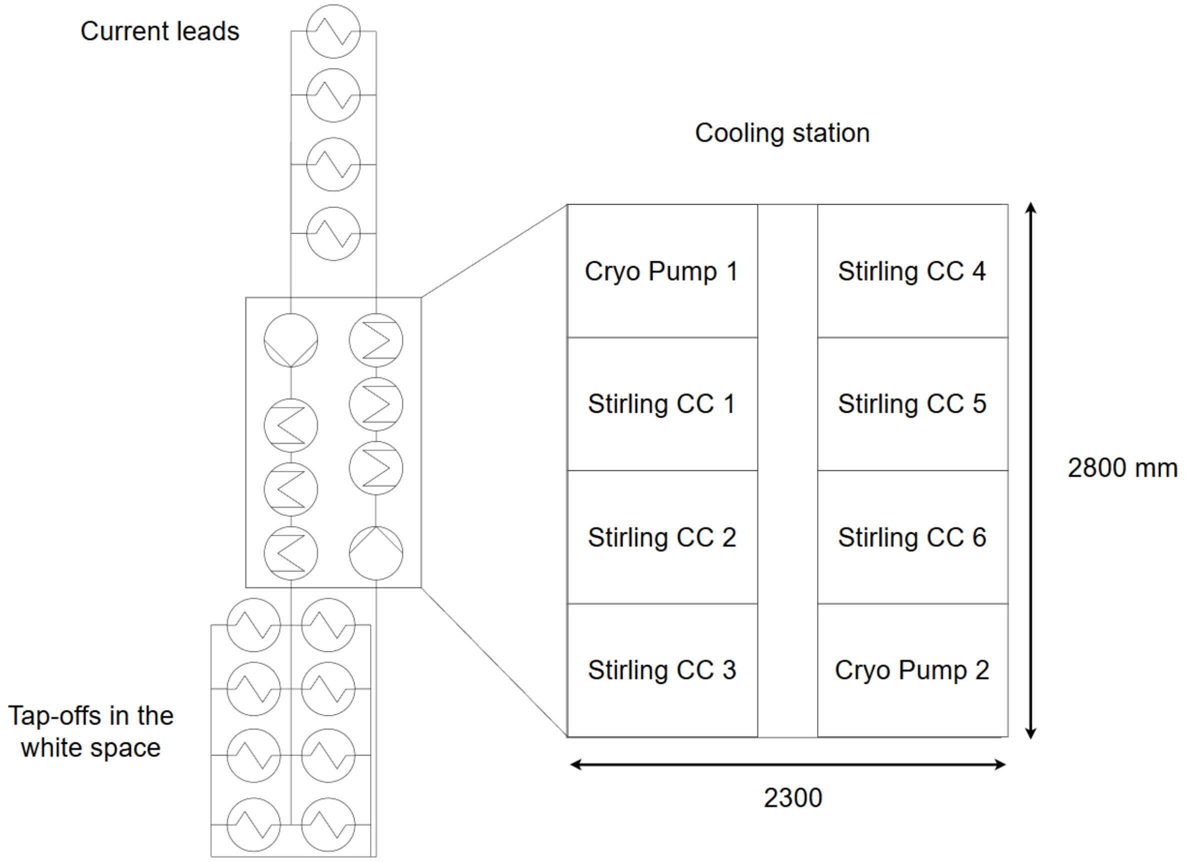

Figure 16 Left: overview of the cooling system and liquid nitrogen flow. Right: top-down view of an exemplary cooling station that can fit inside a container.

2.6. GPU POD

A POD can consist of four rows of 12 GPU racks. Each VR200(Rubin) rated at 200 kW, resulting in total power of 9,600 kW per POD, allowing for additional space for network racks. Each POD is fed by an A and B line to allow for redundancy.

The exact layout of the POD, where the networking racks are in relation to the compute racks and CDUs are still to be decided.

Uninterruptable power supply (UPS) units can be introduced on a POD basis; the main backup power will be supplied by the BESS. Capacitors and high cycle batteries will be used at the rack level to smooth load spikes. They can respond very quickly and with low losses, making them well suited for transient load compensation.

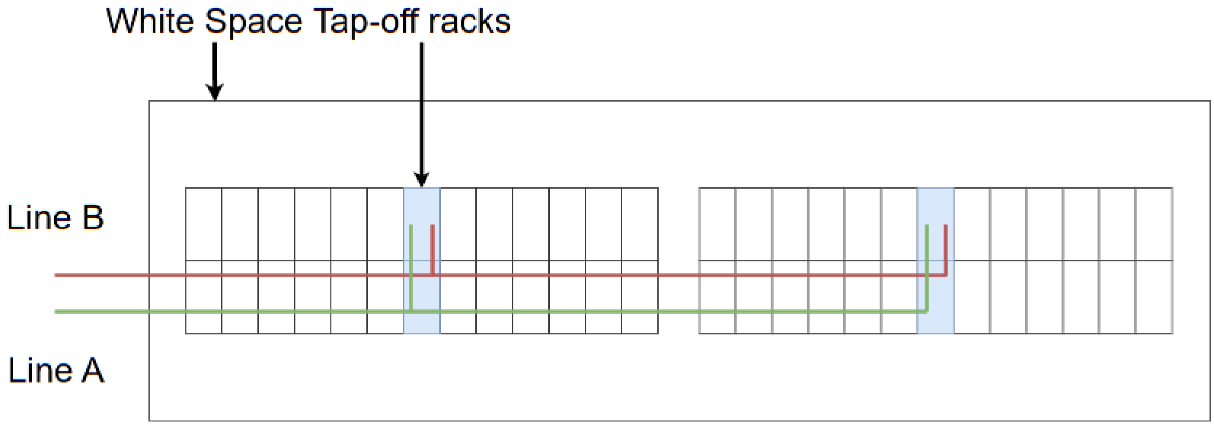

Figure 17 Compute block layout with a tap-off sidecar. Includes 4 rows of 12 GPU racks.

2.7. Auxiliary Power

Auxiliary and facility cooling systems, lighting and other low power workloads will be served by a low voltage AC connection. Two transformers will supply 480 VAC to the cooling stations and HVAC systems.

3. Concluding Remarks

The presented HTS system illustrates how superconductors can be leveraged to enable a modular data center architecture designed for high-power GPU racks. The internal composition of the compute core and POD is not described in detail, as it is not central to the concept. The key design principles are the high current-carrying capability and the use of an 800 VDC architecture, which can be reliably supported by superconducting materials.

The HTS system enables a compact, low-footprint deployment of high-power GPU racks without the need for extensive copper infrastructure or intermediate medium-voltage conversion stages.

All required components of the HTS system are already industrialized and widely available. HTS tape production has increased significantly in recent years, alongside continuous improvements in current-carrying capability and manufacturing quality. Demand for HTS tapes is expected to grow steadily in the near future, driven by applications in power grids and fusion energy systems. Multiple global manufacturers are currently producing HTS tapes with further capacity expansions planned.