Toward a Metal Usage Effectiveness metric in Data Centers: A Study

Marie Reinbigler†
SAMOVAR, Inria Saclay, Télécom SudParis, Institut Polytechnique de Paris
Palaiseau, France

Rosalie Martin†
Adobe Research
Paris, France

***Abstract*—Through a detailed breakdown of the metals required to manufacture and interconnect IT components in a High-Performance Computing (HPC) supercomputer, we introduce a novel metric for evaluating its environmental impact: the Metal Usage Effectiveness (MUE). This metric is based on the abiotic depletion potential (ADP), an indicator of the non-renewable material consumption relative to remaining resources. We define this metric as the ADP (expressed in kg antimony equivalents) caused by the extraction of metals required to manufacture and interconnect IT components, divided by the supercomputer computing power (in FLOPS). The MUE, insightful on the supercomputer manufacturing cost relatively to its computing power, complements existing metrics mainly focused on the use phase cost, such as the Power Usage Effectiveness (PUE) or the Water Usage Effectiveness (WUE). It could be a first step toward a more complete Life Cycle Inventory (LCI) of supercomputers and, by extension, data centers. We provide models for estimating copper and gold content in compute hardware components and interconnects of an HPC supercomputer, and apply them to Frontier, the world's first exascale supercomputer.**

***Index Terms*—HPC, Supercomputer, Data Center, Metal, Sustainability, Metric, LCA, LCI**

I. Introduction

At a time when a new data center is built daily [1], measuring the environmental impact of these constructions is crucial. Over a hundred metrics exist to classify and monitor data centers, covering aspects ranging from energy efficiency, cooling, greenness, performance, network, storage, and security. Among them, the Power Usage Effectiveness (PUE) is the standard to indicate energy performance, like its counterparts for water consumption (Water Usage Effectiveness, WUE) and carbon emissions (Carbon Usage Effectiveness CUE). They measure the proportion of energy/water consumed or carbon emitted by an entire data center relative to the energy consumption of its IT equipment. These indicators are essential to ensure a trade-off between high performance and limited environmental impact, such as access to water and electricity, especially for local populations.

Despite the abundance of metrics, none account for the metals, minerals, and other raw materials required to build data centers, especially the IT equipment. Central Processing Units (CPUs), Graphics Processing Units (GPUs), storage, and cables are all highly demanding of metals for their manufacturing and constitute the core of the IT equipment of data centers and supercomputers. As AI growth causes an unprecedented demand for data centers, the repercussions for the mining industry should be carefully monitored, particularly given increasing metal rarefaction and competing demands from other sectors like energy transition and electric vehicles.

The Life Cycle Assessment (LCA) of a product measures the environmental impact from its manufacture, through its distribution and use, to its recycling or final disposal. It begins by accounting for the extraction of raw materials, commonly measured as abiotic resource depletion and reported in an LCA as Abiotic Depletion Potential. ADP is derived for each extracted elements and fossil fuels and is expressed relatively to the depletion of the element antimony (Sb), representing one gram per ton of Earth's crust. Its unit is the kg Sb equivalents.

A major challenge in inventorying raw materials in IT components, especially high-end GPUs and CPUs, is that manufacturers don't publicly disclose the composition of these products. While the metals commonly found in semiconductors, connectors, Printed Circuit Boards (PCB) circuits, and cables are known, their exact quantities remain undisclosed.

In this study, we present a detailed breakdown of the IT components involved in computations of the Frontier supercomputer, down to the micro level of the components, focusing on copper and gold. Copper is omnipresent in IT hardware thanks to its good conductivity and has been declared in 2023 as critical materials for electricity in the United States [2] and as critical raw materials in the European Union [3]. Gold is a highly-conductive metal used mainly for connectors, present in negligible quantities compared to copper, but five orders of magnitude more expensive and rarer than copper [4], [5]. Our contributions are the following:

- A set of models for estimating copper and gold content at various scales of compute hardware: microscale (in-package inter-chip connections), mesoscale (PCB, connectors, and package-based components) and macroscale (component interconnects and power supply).
- A component breakdown of the computing hardware of a supercomputer, up to the microscale level.
- An estimate of the copper and gold volumes used in the IT components of Frontier supercomputer.
- The formulation of a novel metric, the Metal Usage Effectiveness (MUE), insightful on the raw materials required to construct a data center relative to its computing power.

† Equal contribution

979-8-3315-8717-8/25/$31.00 ©2025 IEEE
DOI 10.1109/ICT4S68164.2025.00024

II. Related Work

Data center metrics. A wide range of metrics is used to evaluate and monitor data centers and supercomputers, encompassing computing power, energy and water usage efficiency, thermal and humidity indicators, reliability, security, network, and storage capacity. [6]–[9] provide definitions and insights on more than a hundred data center metrics.

Most of them are de facto aligned with the data center cost-effectiveness. This is the case for performance metrics, such as FLOPS per Watt (FpW) [10], energy efficiency metrics [11]–[13], such as PUE [14], [15], Space Watts and Performance (SWaP) [13], as well as for cooling system efficiency metrics, such as Water Use Effectiveness (WUE) [16]. Most of them are also used to assess data center sustainability, as optimizing them involves minimizing resource requirements and waste.

The Green Index (TGI), introduced by [17], provides a single score to rank High Performance Computing (HPC) systems: performance-per-watt. This metric is used to rank supercomputers on the Green500 list. [9], [18] further propose classifications and sets of metrics to effectively assess the data centers' sustainability. [19] and [20] propose alternative sustainability scores to replace the performance-per-watt metric.

Only a few metrics are designed to solely inform the sustainability of a data center, independently of cost or waste reduction perspective, such as the Carbon Usage Effectiveness (CUE) [21] related to carbon emissions, and the Green Energy Coefficient (GEC) [22] reflecting the use of renewable energy.

Regarding materials and metals, few metrics on data center components recyclability exist, the Material Recycling Ratio (MRR) and its inverse the Material Reuse Effectiveness (MRE) [23]. However, none of the hundred metrics report on raw material quantities required to assemble data center components.

Life-Cycle Assessment (LCA) and Abiotic Depletion Potential (ADP). As the end-of-life management of data center equipment is a topic of growing interest to reduce the e-waste, many providers are trying to increase the proportion of recycled components, both in the incoming and outgoing flows [24]. This is done by optimizing the Material Circularity Indicator (MCI) [25], a methodology defined for any kind of product that assesses the circularity of a company's product and material flows. The MCI complements the LCA, a process that evaluates the environmental impact of a product or service throughout its life cycle, from cradle (raw material extraction) to grave (recycling or final disposal) [26]–[29].

The first LCA step, the Life Cycle Inventory (LCI), consists of a detailed collection of all the components (up to raw materials) and resources (energy, water) consumed by the analyzed product. [30]–[32] established the formulation of the ADP of a manufactured product.

There is an increasing availability of LCA for ICT products [33], although for the high-end hardware components used in modern data centers and HPC facilities, especially the GPUs, it remains extremely difficult to proceed to the material inventory as the manufacturers do not disclose the composition publicly.

Raw materials in hardware components. The choice of metals used in the IT hardware components is driven by their physical properties [5], [34]. Copper is used in PCBs, cables, and connectors of all the IT hardware components. Gold is mostly used for connectors and pins. [5], [35]–[37] intent to provide a complete study on the average composition of raw materials of all the IT components, hence allowing to assess the share of ICT sector on metals demand. These studies report on average hardware and don't allow to be specific to a particular supercomputer.

Some papers have estimated the metal content of some equipment through the recycling process. For Random Access Memory (RAM) modules, [38] and [39] reported copper content ranging from 19 to 27% of the weight and gold content from 0.10 to 0.13%. Compute storage, particularly Solid-State Drive (SSD) recycling process has been studied by [40]. It is similar to any chip package one and has both copper and gold content [41]. PCB recycling has been extensively studied. [42]–[45] estimated a metal content ranging from 10 to 30% for copper, and from 0.002 to 0.056% for gold. Power supply busbars used in data centers are composed of a 99.9% copper alloy [46], contributing significantly to the total copper weight.

III. Method

In this section, we present our models for estimating the volume of copper and gold in each compute component, from the microscopic to the macroscopic scale.

A. Microscale components

The microscale study concerns all copper/gold-based components present within a chip package. The metallic elements in a package are interconnected via horizontal or vertical connections as illustrated in Fig. 2.

1) Horizontal interconnections: Horizontal interconnects provide inter-chip communications using wires, and the possibility to move an I/O pad access to another location on the chip ReDistribution Layers.

Wires: Wires are used for chip-to-chip interconnections. It is currently performed by traces integrated in the substrate, or for high-performance interconnections in an interposer. Depending on the features of the material it is included in, the pitch and line width have to be adapted. It is characterized by the wire width w and its pitch. Using an estimate of wires length λ and their metal plating thickness ϵ, usually made of Cu, we can estimate the metal volume used in the package. Thus, $V_{Cu} = w * \lambda * \epsilon_{Cu}$.

ReDistribution Layer (RDL): It eases the vertical integration. It is composed of traces having similar properties as wires, meaning width and inter-line spaces. Based on their length and copper plating thickness, the metal volume is determined.

2) Vertical interconnections: Vertical interconnections enable stacking packages. Packages initially integrated chips in a 2D manner, with a chip on top of a substrate, which is on top of a PCB. Thus, It requires solders for substrate-to-PCB integration and chip-to-substrate connections. With the emergence of advanced heterogeneous packaging technologies in 2.5D/3D, additional elements are stacked, thus requiring

Fig. 1: Top-view of Frontier compute area, with the cabinet power supply busways.

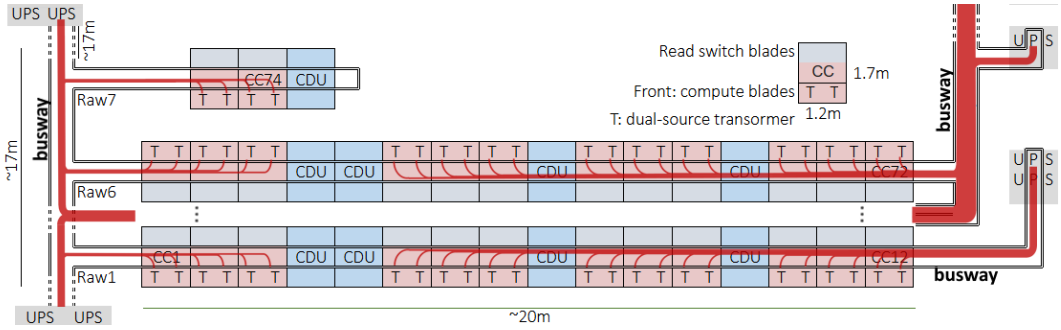

additional vertical technologies like micro bumps, Cu-Pillars, and Through-Silicon Vias (TSV) [47].

Solder and micro bumps: Solders are balls of metal used to bond the printed circuit board to the substrate [47]. The C4 bumps usually link the substrate to the above dies. Microbumps are used for additional stacked chips in multichip packages: they feature smaller diameters and pitch sizes. [48]. Usually, the bumps are composed of a metal alloy, often containing copper with a proportion ϕ_{Cu}. Thus, $V_{Cu} = V_{Ball} * n_{Ball} * \phi_{Cu}$.

Cu-pillars: It consists of a copper pillar with a solder cap. It maintains height while reducing solder diameter [49].

Through-Silicon vias: Metallic tubes, usually made of copper, characterized by both their height h and diameter d, or the aspect ratio [50]. Thus, $V_{Cu} = \pi * (\frac{d}{2})^2 * h$.

B. Mesoscale components

Mesoscale components refers to all components hosting the microscale components, components package, and out-of-package elements like connectors.

1) Printed Circuit Board and Substrate: In high-end computers, PCBs are usually multilayered PCBs. It consists of two outer layers of copper foil and n number of inner layers as displayed in Fig. 3b. Each inner layer is composed of a core layer, glass, or substrate, surrounded by two copper layers, also in between two prepreg layers. Each inner copper layer is trimmed to leave only the desired traces for further interconnections. Fully copper-plated vertical vias are then formed to connect the different PCB layers [51]. If the map is known, as well as the number of vias, their length, and diameter, the copper content can be estimated from the copper volume. Otherwise, it can be estimated by evaluating the copper content of each copper layer without tracing. In addition, at the end of the process, the PCB is plated with gold pads to provide high-quality interconnectivity with the bonded chips. To estimate this value, we need to consider both the thickness of the gold plating and the total number of connections required to wire all the packages to be bonded. For a Multilayered PCB of size $w * h$, the copper volume can be estimated by $V_{Cu} = w * h * (2 * \epsilon_{outer} + (n-2) * \epsilon_{inner})$ and gold volume by $V_{Au} = n_{connection} * S_{pad} * \epsilon_{gold}$.

The substrate has a structure close to the PCB one. It is described by a triplet $(n_{top} - n_{core} - n_{bottom})$ providing the number of metal layers around a set of core layers. If no detail about how layers are interconnected, the methodology will be the same as for the PCB. Otherwise, a more accurate estimation can be performed.

2) Connectors: Various connectors are involved at the scale of a compute node, consisting of a node board with plugged-in components. CPUs connect via sockets, while GPUs, Network Interface Cards (NIC), and storage use dedicated connectors. There are also inter-package network connectors.

Although they can have different shapes, connectors share common features. They have a set of pins or pads whose conduction is achieved by metal plating. They are also two-sided, with a male and a female connector belonging to two different components (Fig. 3a). Based on the thickness, height, and width of the plating, it is possible to estimate the volume of metal involved. In our study, we will consider gold as plating material. Thus, $V_{Au} = n_{pads} * S_{pad} * \epsilon_{gold}$

3) Packages: The metal content of the following components is a combination of microscale components, PCBs and/or substrate embedding various components. Thus, we detail each component specificities, with metal volumes estimated using previous formulas.

GPU: Based on the NASA report [52], a GPU board is mainly composed of one or several Graphical Compute Dies (GCD) with efficient in-package interconnections, linked to a set of High-Bandwidth Memory (HBM) stacks, a set of capacitors and, are on top of a at least a substrate, itself above the GPU board. Among those components, copper and/or gold are present in the GPU board, a PCB of up to 8 layers [53], in its connector to the motherboard, in the in-package inter-chip connections, and in the vertical integration components. The capacitors can contain metals, in general palladium or aluminum [54] and won't be considered in the context of this study. The HBM stack is composed of several memory modules on top of a memory controller. Each memory module is interconnected to the rest of the stack via a set of TSV and microbumps. The whole stack is linked to the rest of the package and to the GCD. The number of each category of interconnections is dependent on the version of the HBM and is described by the corresponding datasheet.

Fig. 2: Multi-Chip Module (MCM) packaging technologies. 2D refers to chips on top of a substrate that enables interconnections. A 2.5D package contains an additional interconnection layer, an interposer or a bridge, with higher density. It can also contain chip stacking. The 3D technology stacks massively the chip to reduce interconnect length and increase its density. [47]

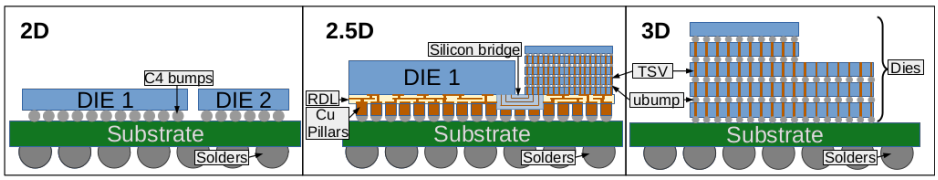

In addition, the GPU package can be characterized by two metrics: the areal interconnect density, *i.e.* the number of vertical interconnects per unit area, and the linear interconnect density, *i.e.* the number of horizontal interconnects [55]. Using these two metrics, we can estimate the number of horizontal interconnections based on the GCD die perimeter and the linear interconnect density. Using the substrate or interposer features, we can access the width of each wire. With the plating thickness and the average wire length, we can get an approximation of the copper weight involved in linear connections. For the vertical interconnections, depending on the GPU package architecture, we know which kind of interconnect is used at which layer, as well as their shape characteristics. By having an idea of the proportion of interconnects used per layer, we can estimate the volume of copper involved.

Processors: In our model, we consider a processor as a Multi-Chip-Module (MCM) as high-end processors would require increasing die size to continue performance improvement which is not realistic from a financial point of view [55]. Thus, the processor package contains a base substrate on top of which we can find a set of CPU units, with or without an I/O card depending on the architecture. The package is connected to the motherboard via a socket. Each chip is connected to the substrate via a set of vertical connectors, and, depending on the format of the package, to other chips via linear interconnects.

Node memory: High-performance flash memory needs to be provided to each compute node. Thus, the most adopted technology is the SSD [55]. There are 3 main form-factors categories with their dedicated connectors: the Enterprise Datacenter Standard Form Factor (EDSFF), the M2 form factor, and the U2 form factor [56]. The last one is progressively being abandoned due to its larger size. The two others are available in a wide range of sizes. EDSFF is based on the efficient NVMe protocol with PCIe interfaces. M2 form factor can support NVMe protocol with PCIe interfaces or SATA protocol and interfaces. Using these specifications in combination with the SSD dimensions, we can estimate the copper and gold content of the PCB as well as the gold content of the connectors.

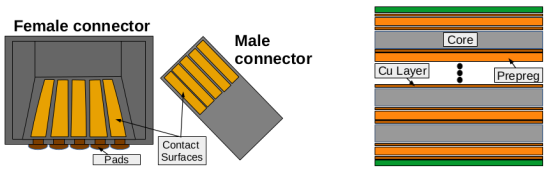

(a) Connector model (b) PCB model

Fig. 3: Mesoscale components

Network Interconnect Card: The NIC is a package containing a substrate embedding two NIC cores. It makes the link between the computer and the network computer via a network connector, *e.g.* an ethernet port, as well as to the motherboard, using *e.g.* PCIe. Based on the PCB and substrate dimensions, their number of layers, and connector specifications, we can estimate the copper and gold content of the NIC component.

Network Switch: Similar to the NIC, the network switch is a PCB with a substrate that houses the network compute core. It provides the interconnection between the network ports of the NICs. It thus contains a set of components making the interconnections between all the ports via the substrate, with additional function blocks to manage it and to perform network algorithms such as Ethernet lookup or maintaining the routing table. The main sources of copper and gold are in the PCB and substrate layers, as well as in all network connectors.

Random Access Memory: The CPU requires access to RAM modules via the motherboard and a Dual Inline Memory Module (DIMM) port. The DDR RAM generation provides the module shape and the DIMM connector specifications. Copper is mainly present in the RAM PCB, while gold is omnipresent in the DIMM pins of the card, its port, and PCB pads.

C. Macroscale components

This section is mostly dedicated to the interconnect and power supply of components in a data center by means of cables. As a general rule, estimating the volume in cm^3 of the conductor metal used in a cable is given by:

$$V_{cable} = \kappa * \sigma * \lambda * 10^{-2} \quad (1)$$

where κ is the number of wires composing the cable, σ is the cross section of each wire, in mm^2, and λ is the length of the cable, in cm. The cable section is in general provided in AWG (American Wire Gauge, where a larger section area has a lower AWG), which we convert into mm^2. We consider only the wire's volume in this section, not their connectors.

1) Node and switch interconnect - HPE Slingshot: In our study focusing on large-scale computing systems, we consider the most powerful interconnection network system at the time

of this paper: HPE Slingshot 11 [57]. It allows building exascale and hyper-scale data center networks with at most three switch-to-switch hops, with nodes connected in a Dragonfly topology [58]. In this topology, compute nodes are regrouped inside a compute group, and each compute node is connected to x_{L0} switches inside the group (L0 links). All switches inside a compute group are fully interconnected (L1 links) using x_{L1} links, and each compute group is also fully connected to other compute groups (L2 links), using x_{L2} links.

L0: In group Switch-to-node: Copper cables are widely used for this local connection from the switch to the compute node, either with Direct Attached Cables (DAC) with QSFP-DD connectors or with a cable assembly ending with ExaMAX connectors, reaching directly the compute node's NICs. Whatever the format, those cables are made of n pairs of copper wires with a small section σ and length λ lower than 150cm.

L1: In group switch-to-switch: DAC copper cables with QSFP-DD connectors are the first choice for these local connections. The amount of L1 links to fully interconnect all switches of a compute group equals $n * (n-1)/2$, and their length can be approximated given the dimension of the switches, and their arrangement in the compute group.

L2: Other group switch-to-switch: Because copper cables are prone to signal degradation over long distances, optical fibres are generally used for L2 links that connect two compute groups. These are essentially made of silicon and a bit of germanium, which is outside our focus for this study.

2) Power supply wires: To supply power to the components in an HPC supercomputer, the power flows from the utility power to an Automatic Transfer Switch (ATS) allowing to switch to a backup power generator in case of interruption, before reaching a set of Uninterruptible Power Supplies (UPS), installed close to the compute cabinets. The UPS supplies the compute cabinets with wires installed with a busway system, with redundancy ensured by dual-source transformers. The busway is generally mounted overhead and the transformers on top of the rack they supply, below the busway.

A dual-source transformer supplies the in-rack Power Distribution Units (PDU), which provide circuit breakers and further supply a set of Power Supply Units (PSU), that power the compute nodes with a constant energy. Inside the node, the power is stepped down by components like a Super Intermediate Voltage Converter (SIVOC) to supply the node components (CPU, GPU) [59], [60].

Copper wires are used for all connections, with decreasing capacity through the flow, that we will account for from the Uninterruptible Power Supply (UPS) cabinet to the end components in a compute node.

UPS to dual-source transformers: Each compute cabinet is supplied with one or several dual-source transformers. Each power supply wire from a UPS to a transformer follows the busway system below the ceiling of the Compute Room. For each dual-source transformer, there are two 3-phase copper wires of a big cross-section σ from one UPS. The length λ of the wire equals the path length using the busway between the UPS cabinet and the transformer, located on top of the compute cabinet. It could be from a few meters to hundreds of meters. Equation 1 applies with κ equals 3.

Dual-source transformer to compute cabinet: Similarly, for each dual-source transformer, two 3-phase copper cables reach the cabinet directly below ($\lambda \leq 5m$). These cables can then supply an AC Busbar, to which the PDUs are connected.

Compute cabinet PDU to PSU: Each in-rack PDU can supply x_{PSU} PSUs in the cabinet, still with 3-phases copper wires, so equation 1 still applies, but with a smaller section σ. The power in the output of the PSU is in general single-phase DC and may step down the voltage.

PSU to compute node: A PSU can supply compute nodes and other components in a compute rack by means of a DC busbar. Single-phase copper wires are used for this connection, where each PSU connects to the DC bus, so we count x_{PSU} wires per cabinet, for which equation 1 applies with κ equals 1. The DC bus further connects to compute nodes directly with connectors to the power input of the nodes.

Compute node to IVOC: The power in input to the node further supplies the Intermediate Voltage Converters (IVOC) of the node, which again is a power supply that converts input high voltage to the appropriate low voltage that can be used by the various components within the node. Cables assembly with κ wires of small section σ is used here, allowing to split the amount of cables from the input into several groups of cables to reach the IVOCs.

3) Busbars: A busbar is a bar of solid, made of conductive metal, used to transport an electric current with minimal energy loss. In a high-end supercomputer busbars are composed of copper for its highest level of conductivity, they can be laminated with layers of copper separated by layers of insulator, or not. They also minimize the amount of cabling in the limited space of the compute chassis.

The UPS and the compute cabinet contain one or several busbars, that can handle single or three-phase current. As it is roughly a solid of metal, the volume of metal can directly be estimated with the dimensions of the busbars.

4) Compute node & blade: A compute node consists of a node board (PCB) with CPU(s), GPU(s), NIC(s), RAM, and local storage connected to it with connectors or through the PCB. We can aggregate the volume of copper and gold at the compute node level by estimating the volume of the node board and adding it to the embedded components.

A compute blade hosts one or more compute nodes in a metallic assembly, further contained within a compute chassis. It presents a backplane with connectors, such as ExaMAX, to communicate with the other elements of the chassis.

5) Compute cabinet: A compute cabinet is a structure that hosts nodes and switches mounted in chassis, with PSUs, PDUs, busbars, and power cables, and network interconnection within the cabinet and to other compute cabinets. These elements are arranged in a super-efficient manner to maximize power efficiency and minimize space and cable waste.

Newer data centers and supercomputers also use a water cooling system that is fully integrated across all IT components. A compute cabinet also contains a management compo-

nent that controls all of its components and an environmental controller. We ignore these parts in this component breakdown, but they should be considered in future work.

IV. Results

In this section, we apply our models to measure the MUE of Frontier supercomputer, which is, at the time of the writing, ranked second in the Top500, listing the most powerful computer in the world in terms of FLOPS, and was the first worldwide to reach exascale computing. [61]–[63] provide a precise description of the node design, the interconnect, the I/O subsystem, the power supply, the software environment, and the cooling system. [57] describes the Slingshot interconnect technology, deployed in a dragonfly topology [58].

TABLE I: IT hardware components used in Frontier and their count per blade, per cabinet and for the 74 cabinets of Frontier.

Comp.	Name	Blade	Cab.	Frontier
PDU	100A w/ 4circuit breakers		4	296
PSU	15kW Rectifier		32	2368
Switch	HPE Slingshot Switch		32	2368
Blade	Cray EX 235a	1	64	4736
Node	Bard Peak	2	128	9472
CPU	3rd Gen AMD EPYC Trento	2	128	9472
GPU	AMD Instinct MI250X	8	512	37888
GCD	AMD Instinct MI250X GCD	16	1024	75776
NVMe	M.2 NVMe 1.92TB	4	256	18944
NIC	HPE Slingshot NIC	8	512	37888
RAM	64GB DIMMs of DDR4-3200	16	1024	75776
SIVOC	HPE SIVOC 3200W 380V-48V	2	128	9472

TABLE II: Frontier components dimensions, captured from [64]–[68] and extrapolated from pictures of Cray EX 235a.

Component	Width (cm)	Height (cm)	Depth (cm)
Cabinet	120.00	200.00	170.00
Switch blade	48.00	3.81	30.00
Compute Blade	100.00	3.81	44.45
Node card	40.45	37.63	-
CPU	5.85	7.54	-
GPU	26.70	11.10	-
NIC	8.91	18.82	-
SIVOC	4.61	16.89	-
RAM	13.30	3.10	-
NVMe	8.92	21.40	-

We restrain our study to the compute part of Frontier, made of 74 HPE Cray EX4000 compute cabinets [64], [65], [69], each cabinet being composed of 64 HPE Cray EX235a blades and 32 fully interconnected HPE Slingshot switch blades [70], assembled within 8 compute chassis (front door) and 8 switch chassis (rear door) in the cabinet. The HPE Cray EX235a blade contains two compute nodes, named Bard Peak, made of one 3rd Gen AMD EPYC™ Trento CPU and four AMD Instinct MI250X GPUs (e.g. 8 Graphic Compute Dies), two 2TB NVMe, and four Network Interface Cards (NIC). Table I summarizes the main components of the compute cabinet of Frontier, and Fig. 4 illustrates how compute blades and switch blades are arranged orthogonally and connected directly through ExaMAX orthogonal connectors. We also gather the measures of these components in Table II.

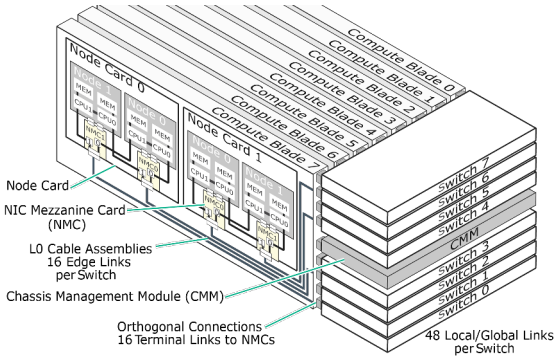

Fig. 4: Orthogonal assembling of compute blades and switch blades in their respective chassis (not shown). The compute blade presented differs from HPE Cray EX235a. Source: [69]

We ignore the cooling system hardware components as well as the storage cabinets related hardware. The power supply breakdown is limited to the Frontier Compute Room perimeter. All these aspects may be added in future work.

A. Metals in connectors

TABLE III: Connector features and their volume of gold and copper. Information from [66], [68], [71]–[76].

Type	# Contacts		Surface (mm^2)		Volume (mm^3)	
	Total	to PCB	Male	Female	Au (M/F)	Cu
QSFP-DD	162	152	176.84	176.84	0.067	-
ExaMAX	112	112	464.00	464.00	0.35	3.02
PCIe	328	328	493.64	218.12	0.38 / 0.17	-
RDIMM	288	288	414.00	453.00	0.31 / 0.34	-
OAM	688	688	860.00	270.18	0.65	-
SP3	4094	4094	2338.09	803.86	1.77 / 0.61	-
M2	75	75	84.38	67.50	0.064 / 0.051	-

IT components embed connectors. Thus, we first compute their metal composition per type based on their features (Table III) before combining them with microscale components to form mesoscale components. In this subsection, gold plating thickness is set to $\epsilon = 0.76\mu$m.

QSFP-DD connectors: This connector contains gold in its connexions pads, both from the male and female sizes. Thus, considering the pad surface and the plating thickness, each QSFP-DD contains 6.72×10^{-2} mm^3.

ExaMAX connectors: The ExaMAX connector features curved pins with gold-plated contact surfaces, made from a copper alloy. Frontier uses the 4-pair/8 column ExaMAX connector. Based on the plating thickness and the pin's external surface, the gold volume is 35.2×10^{-2} mm^3. Considering a pin volume of 3.36 mm^3 [73] and assuming the alloy to be 90%Cu, each connector contains 3.02 mm^3 of copper.

PCIe connectors: The PCIe 4x16 connector is asymmetrical between its male and female versions. In both cases, the connection is performed via gold-plated pads distributed on two rows [74]. However, their surface differs: 493.64 mm^2 for the male connector and 218.12 mm^2 for the female one, *i.e.* 0.38 and 0.17 mm^3 volume of gold respectively.

RDIMM connectors: The RDIMM connector for DDR4 RAM has similar features than PCIe 4x16 connector: the gold-plated pads surface is 414 mm^2 and 453 mm^2 for male and female connector, *i.e.* 0.31 and 0.34 mm^3 gold volume [68].

Open Accelerator Module (OAM) connector: The OAM connects the AMD Instinct MI250X GPU to the motherboard. There are two connectors per GPU board. The male and female connectors are identical. They are composed of gold-plated 688 pins, *i.e.* 0.65 mm^3 of gold [71].

Socket SP3: The SP3 socket is the connector of the CPU to the motherboard. The connection is performed via pads on the CPU side and a set of pins from the socket size [66]. The pad surface is larger than the socket pin surface. Thus, their gold volume differs: 1,77 and 0.61 mm^3.

M2 connector: The SSD disks are connected to the compute node via M2 connectors. The contact is performed via gold-plated pads, 84.2 mm^2 for the male connector and 67.5 mm^2, *i.e.* 6.41×10^{-2} and 5.13×10^{-2} mm^3 [75], [76].

B. Components in Frontier

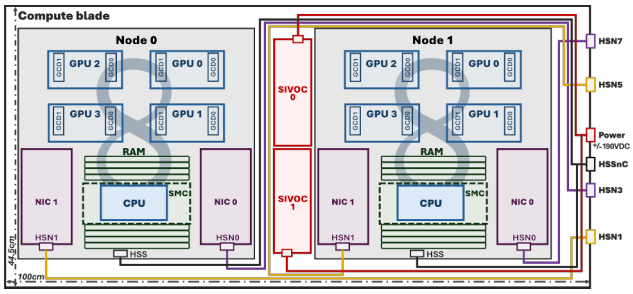

Fig. 5: Cray EX235a diagram. Components are represented with their actual scale and cables follow the correct path. Each HSN link is a dual Slingshot port (two 200 Gbps NICs). The power supply links between the SIVOCs and the components are not shown for clarity.

In this section, solders are assumed to be made of a copper alloy with 0.5% of copper. The gold-plated pads surface on top of high-performance PCBs connections is assumed to be 1 mm^2. The PCB external copper layers are set to 0.0175 mm thick and the inner layer to 0.035 mm thick according to standard PCB structures [83]. The gold and copper volume accumulation are in Table VI. The copper is assumed to occupy 50% of PCB layers, except the node motherboard where 30% has been selected as way more layers needs to be stacked.

CPU: The CPU is a 2D package composed of 5 main chips, 4 complex core dies and one I/O die, on top of a substrate of 18 metal layers. The metal is assumed to be copper, and following our model, we overestimate copper content by considering 18 copper layers like for EPYC™ Rome CPU model [66], with a thickness of 0.0175 mm, similar to standard PCB copper foil thickness. The horizontal interconnectivity between dies is performed by InfinityBand lanes contained in the substrate. The vertical connectivity is performed by C4 bumps. The whole package follows the SP3 package Land Grid Array format of 4094 pads. Based on [84], we consider for C4 bumps a diameter of 100 μm with 250 μm pitch.

GPU: The GPU is a 2.5D package composed of 10 main chips, 8 HBM stacks dies, and 2 Graphics Compute Dies (GDC). The GCDs are interconnected via InfinityFabric lanes and to HBM stacks via Silicon Bridges for local higher interconnect densities. It also contains a Redistribution Layer with a coarser interconnect density with a 2 μ line width and space. From the pictures, we considered 6 metal layers inside the substrate. The metal is assumed to be copper, and following our model, we estimate copper content by considering 18 copper layers of 0.0175 mm, similar to standard PCB copper foil thickness. We considered for the RDL layer and the silicon bridges that all interconnections have the same length with a distance equal to the core pitch assumed to be 100 μm [85]. The silicon bridges are located between HBM stacks and GCDs and inter-GCDs. We assume that the bridge is present along the edge of the smallest die, with 11mm width for connections to HBM and 25mm between the dies. Considering a linear interconnect density of 750 links/mm/layer [86], there are 84750 links in the silicon substrate. With a GCD perimeter of 111mm, there are 63000 additional links in the RDL layer. In total, with a copper plating thickness of 0.0175 mm, these wires account for 0.2798 mm^3. The areal interconnect density provides the bump quantity per mm^2. The applied value here is 650 [86]. Knowing a HBM2e die has 6303 C4 bumps and 5000 TSV of the same number of micro-bumps, we end up with 5000 TSV of length 0.72 mm and 0.025 mm for the diameter. The TSVs are made out of copper, i.e. 1.77 mm^3. C4 bumps are assumed to have the same size as in the CPU with only 0.5% of copper. The connection between the substrate and the GPU board is expected to be composed of solder balls. The solder balls' contribution is computer considering of 0.635mm diameter and 0.535 mm height and a density of 1.8 balls per mm^3. The GPU board PCB is estimated to have 8 copper layers, with two external layers of 0.0175 mm thick and an internal layer of 0.035 mm thick. The gold content is essentially included in the 2 OAM connectors connecting the GPU to the motherboard and the PCB pads at bond points.

NIC: The Network Interface Card is a mezzanine card in HPE Cray EX235a [69]. It is composed of two Network Core dies on top of a substrate. The interconnections in the substrate are in this case ensured by 1854 TSVs. Using [87], we assumed their diameter to be 0.125 mm and their height to 0.6 mm. It connects the motherboard to the blade connector via two PCie 4x16 ports and 1 network port. As the network port type is not provided, we ignore it in this study. The solders considered are the same as for the GPU with the same density. Concerning gold, it is contained in the QSFP-DD connector, in the two PCIe connectors, and in the PCB pads. The PCB is estimated to have 8 layers due to its routing role [88]. One specificity of the mezzanine card is the inclusion of a riser, here being the link between two PCIe female ports. Thus, there are 2 PCIe female and 4 extra PCIe male connectors.

Slingshot Switch: The Slingshot Switch PCB contains 8

TABLE IV: Component Board Features [51], [66], [68], [77]–[82]

Comp.	# PCB Layers	Substrate surface	Substrate Layers	# Solders	# C4 bumps	# TSV	# ubumps	RDL (#/L/S)	# Dies	Σ Dies surface	Connectors
CPU	-	4411	9-2-9	4094	393	-	-	-	5	712	1 LSA-4094
GPU	8	9456	3-1-3	17115	521024	5000	5000	1/2/2 1/0.4/0.4 (Si)	10	2328	2 OAM
NIC	8	1024	-	1854	-	1854	-	-	2	264	1 NET, 2 PCIe 4.0x16
RAM	4	-	-	1404	-	-	-	-	18	-	1 RDIMM
NVME	8	-	-	6369	-	-	-	-	16+	3388	2 M2, 1 PCIe 4.0x16
Switch	8	3906	-	3709	-	3709	-	-	64	3516	8 ExaMAX, 24 QSFP-DD

layers [87]. The connection from PCB to substrate is ensured via solders with similar properties as previously and based on a similar switch structure [89]. The copper is also present in the ExaMAX connectors pin volume. The gold is located in all female connectors as well as on the PCB contact points.

RAM: A RAM module is a 4-layer PCB [90] on which 16 DRAM dies and 2 controller dies are soldered. The solders are 0.37 mm wide and 0.48 mm high. Each module has 78 solder balls [82]. The gold comes from PCB pads and RDIMM male connectors.

SSD: The compute node has access to two 1.92 TB NVMe SSD disks with an M2 connector. We selected the M2 22110 form factor for this study [75], [76]. Two of them are included in the package of a storage mezzanine card. Thus, we suppose there are also 2 M2 female connectors. The card also contains one PCIe 4x16 port. Due to the need for a riser, we required 2 extra male PCIe 4x16 connectors.

Node motherboard: The motherboard is a thick PCB due to the routing complexity required to manage and interconnect all the node components. According to [79], for high-end applications, a motherboard can have 28 layers or more in its composition. We decided to stick to this value. Concerning gold content, it comes mostly from female connectors and gold-plated pads at each bond location. In total, it contains 8 OAM connectors, 1 CPU socket, 5 PCIe 4x16 ports, and 8 RDIMM slots. In addition, they all cumulate 13542 points of connections corresponding to a gold-plated pad area.

For reference, assuming a RAM module weighs around 30 g, the weight of copper and gold represents respectively 13% and 0.08% of the total weight. Based on [38], [39] recovered percentages (See Section II) from older generation RAM modules, we can notice our results are in the right order of magnitude. In addition, for PCB copper content, we simulate component PCBs weight using [91] with the copper layer properties defined at the beginning of this section. From PCB surface and copper layer properties, we obtain 166.12 g of copper for the switch PCB, 1.92 g for the RAM PCB, 22.29 for the NIC PCB, 21.00 g for the NVMe PCB and 46.18 g for the GPU PCB (See Table VI), *i.e.* 30.82%, 10.82%, 27.04%, 22.33% and 31.63% of the PCB weight, which belongs to the expected range proposed by [42]–[45].

C. Metals in cables and busbars

We are mainly looking for copper in cables, using equation 1 to compute the volume of copper using the number of wires κ in the cable, their cross-section σ, and the cable length λ. The copper volume of cables and busbars is reported in Table V. The copper and gold volumes in the ExaMAX and QSFP-DD cable connectors are included in Table VI.

1) Network interconnect: Frontier implements Slingshot interconnect with Dragonfly topology, using an HPE Cray EX class 5 topology with dual injection [70]. A class 5 topology means that all the switches in the compute cabinet belong to the same compute group, so they need to be connected all-to-all (L1 links). Dual-injection refers to the ability to provide 2 Slingshot ports for each compute node (e.g. $x_{L0} = 2$).

L0 connections: The switch and compute blades are directly connected with orthogonal ExaMAX connectors at their backplane (see Fig.4). It removes the need for cables between them and the L0 cables refer to the cable assembly inside a compute blade, that connects each of the 4 ExaMAX connectors to one node's NIC. Four High-Speed Network (HSN) cables made of 16 wires of 30AWG section are used, and we estimate their length using the node diagram that reproduces the path of each cable (See Fig.5) and the component dimensions (Table II). The total L0 cable length in one blade is around 3m80, leading to 18km for 64 compute blades within 74 cabinets, and 9472 ExaMAX connectors.

L1 connections: Following our model, $n * (n-1)/2$ L1 connections are needed to fully interconnect the 32 switches in the cabinet, leading to 496 L1 connections, and each switch uses 31 ports to connect to the others. The switch blade front connectors offer 24 QSFP-DD connectors, e.g. 48 ports to connect to other switches (L1 or L2). To connect all-to-all the 32 switches, 16 QSFP-DD connectors are dedicated for L1 connections, and it is necessary to use bifurcated cables (H shape), for which a connection is 1 port of a 4 ports bifurcated cable, hence $x_{L1} = 1$. It leads to 248 copper bifurcated cables per cabinet, each made of two links of 16 wires of a 28-30 AWG section, depending on the cable length (28 AWG for a $\lambda > 2$ m). To estimate the cable length, we approximate the length of the path between each pair of switches, and couple it with the available Slingshot L1 cable length (0.8 m, 1m,..., 2.6 m), which sums to 553 m of cables smaller than 2 m, and 160m of cables longer than 2 m, for one cabinet. There are also a total 36704 QSFP-DD connectors.

L2 connections: There remain 8 QSFP-DD connectors available on the switch blade for connections to other compute groups and to storage and management groups. Frontier uses optical fiber cables for L2 connections and dedicates $x_{L2} = 2$

links for interconnecting all-to-all the compute groups, using 5 of the 8 L2 connectors for it. We add in our aggregate the gold amount of the 10804 QSFP-DD connectors for L2 cables.

2) Power supply: **UPS busbars:** Each UPS cabinet contains 3 AC copper busbars, which we estimate to be 700 mm^2 section and 1.5 m long, using pictures and power requirements.

UPS to dual-source transformers: Two dual-source transformers are located above each Frontier compute cabinet for redundancy, providing four links per cabinet for a total of 296 links. The Frontier compute area is approximately 340 m^2 and contains 6 rows of 20 m, each with 12 cabinets and 4 CDUs, and an extra raw with only two cabinets and one CDU. The UPS is located against the walls of the Compute Room, which is about 3.5 times the size of the computing area. A pair of UPS provides 48 copper power supply wires, 3-phase with a section of 0AWG each, grouped in the ceiling busway to reach the dual-source transformers and the CDUs (2 links per CDU) of a row. From the many images available of Frontier, we reproduce the map of Frontier in Fig.1 and find that the flow from the four left cabinets goes to the left and continues along a vertical busway to reach the UPS on the bottom and top walls of the room, and the flow from the twelve right cabinets goes to the right through the UPS on the right wall. We estimate the total length of all these cables to be 5.07 km, taking into account the distance between each compute cabinet and the UPS serving its dual-source transformers, including the vertical distance to reach the busways.

Dual-source transformer to compute cabinet: We count four 4 m 0 AWG 3-phases copper cables for supplying one compute cabinet from its two dual-source transformers, making a total of 296 cables and 1.18 km.

Compute cabinet AC busbar: This busbar receives the power supply from the dual-source transformer and is made of 3 laminated copper busbars, that we estimate with a cross-section of 560 mm^2 and 1m50 height, from [64].

AC busbar to PDU: Each of the four PDUs of the compute cabinet is supplied by plugging three 3-phase cables to the AC busbar, with a section of 3 AWG per phase to support 100 A. e estimate their length to 1 m, leading to 0.8 km in total.

PDU to Rectifier shelf: As shown in [64], connectors go from the rear of the PDU to reach the rectifier shelf with the PSUs. There are 24 wires in one cable assembly to reach the eight PSUs of the rectifier shelf, with a section that we estimate to 12 AWG and a 1 m length.

Rectifier shelf to DC Busbar: In the output of the PSUs the power is single-phase, and one rectifier shelf supplies two horizontal compute chassis, with 4 cables (two positives and two negatives) reaching the horizontal DC busbar of each compute chassis. We estimate the section to 8 AWG, and a cable length of 0.8 m, leading to 32 cables per cabinet and 2368 in total, with up to 4.68 km.

Compute chassis DC Busbars: There is one horizontal and one vertical High-voltage DC busbar per compute chassis. The horizontal one is supplied by the rectifier shelf and supplies the compute blades with orthogonal connections, and the vertical one is plugged into the horizontal one, and further supplies the switch blades. We estimate the section of these busbars to 304 mm^2 and their length to be 45 cm.

TABLE V: Specifications and Copper Volume of Frontier Cables and Busbars. The UPS to dual-source transformer power supplies are responsible for the largest copper footprint.

Cables	**Wires**		**Section**		**Len.**	**Vol.**
	κ	#Cables	AWG	σ(mm^2)	λ(km)	V(m^3)
1) Slingshot						**0.0633**
L0	16	18944	30	0.051	18.00	0.0147
L1 $<$ 2m	16	31376	30	0.051	40.91	0.0333
L1 $\geq$ 2m	16	5328	28	0.081	11.84	0.0154
2) Power						**1.3238**
UPS⇒dual	3	296	0	53.47	5.07	1.0197
Dual⇒ACBus	3	296	0	53.47	1.18	0.1899
ACBus⇒PDU	3	888	3	26.67	0.88	0.0710
PDU⇒PSU	24	296	12	3.31	0.30	0.0235
PSU⇒DCBus	1	2368	8	8.37	1.89	0.0159
Node⇒IVOC	16	4736	28	0.05	4.68	0.0024
3) Busbars						**0.4052**
UPS Busbar		54	-	700.0	0.08	0.0567
AC busbar		222	-	560.0	0.33	0.1865
DC busbars		1184	-	304.0	0.53	0.1620

D. Metals in Frontier

Thanks to this precise analysis of the volume of copper and gold in the IT hardware component of Frontier, we can now aggregate the results and get their correspondence in the weight of metal. Given the density ρ of a metal (in kg/m^3) and its volume V (in m^3), the weight in kg of a metal can be computed with $W = V * \rho$. Knowing that at room temperature, copper has a density of 8960 kg/m^3, and gold 19283 kg/m^3, we can deduce the weight of copper and gold for one node, and for the entire Frontier supercomputer.

To express quantities of metals relatively to their impact, we use the ADP, typically used in an LCA of a manufactured product, defined based on [30] by:

$$ADP_i = \frac{DR_i}{{R_i}^2} / \frac{DR_{Sb}}{{R_{Sb}}^2} \quad (2)$$

where DR is the depletion rate (e.g. annual extraction rate), R the reserve of this mineral, i is the observed mineral, and Sb is the antimony element. The ADP of the product is then given by $\sum_i ADP_i * W_i$, in kq Sb eq..

We present in Table VI the metal weight of the IT components, separately and aggregated for Frontier, and provide their ADP in kg Sb eq., using metal statistics from [5].

As a reference, Ademe [5] studied a server with 2 CPUs, SSDs, and Hard Disk Drives, estimating 0.96 g of gold and 1500 g of copper per server. Our node metal weight is of the same order of magnitude. Our higher gold content likely comes from the higher-end Frontier node composition, including 4 GPUs. The lower copper weight may result from excluding the cooling circuit in our estimates.

E. Metal Usage Effectiveness

We define the Metal Usage Effectiveness by the ADP of the IT hardware components, divided by the computing power P of the supercomputer or data center, in PFLOPS:

TABLE VI: Copper and Gold Volume V, Weight W and ADP per component unit and for all Frontier (taking the quantities in table I). Aggregation for the compute components, the slingshot switches and connections and the power supply.

Comp.	**Au** (mm^3, mg)		**Cu** (cm^3, g)		**ADP** (kg Sb eq.)	
	V	W	V	W	**Unit**	**Total**
NodeMB	19.20	370	46.88	420.06	$1.28e^{-2}$	$1.22e^{2}$
CPU	1.78	34	1.39	12.45	$1.17e^{-3}$	$1.11e^{1}$
GPU	14.13	272	5.15	46.18	$9.24e^{-3}$	$3.50e^{2}$
NVMe	1.17	23	2.34	21.00	$7.79e^{-4}$	$1.48e^{1}$
NIC	3.59	69	2.49	22.29	$2.36e^{-3}$	$8.93e^{1}$
RAM	1.35	26	0.21	1.92	$8.81e^{-4}$	$6.67e^{1}$
Tot. Node	105	$2.02e^{3}$	85.24	764	$6.88e^{-2}$	$6.54e^{2}$
Switch Ss	6.41	124	18.54	166.12	$4.31e^{-3}$	$1.02e^{1}$
Frontier	**Au** $W(kg)$		**Cu** $W(kg)$		**ADP**($kg\ Sb\ eq.$)	
Compute[1]		19.178		7234.004	$6.54e^{2}$	
Power[2]		-		15491.385	$1.27e^{1}$	
Slingshot[3]		0.561		960.937	$1.97e^{1}$	
Total		**19.739**		**23686.326**	**6.86**e^{2}	

Compute[1]: Compute blades. **Power**[2]: Power supply cables and busbars. **Slingshot**[3]: Slingshot switches, cables and connectors.

$$MUE = \frac{ADP}{P} \tag{3}$$

MUE should be minimized in order to limit the use of minerals and possibly critical materials while increasing performance. Considering only copper and gold, the **MUE of Frontier is 0.343 kg Sb eq. / PFLOPS**. Table VII summarizes the main metrics of Frontier at the time of this paper.

TABLE VII: Frontier metrics [61]

Footprint	**Power**	**Peak**	**PUE**	**WUE**	**MUE**
370 m2	29MW	2.0 EF	1.03	0.06 [92]	**0.343**

This metric could complement PUE and WUE and provide insight into the balance between power and water efficiency metal usage. For example, liquid cooling has significantly reduced PUE compared to air cooling technologies, but it has increased both WUE and MUE because the cold plate cooling associated with liquid cooling uses a large amount of metal.

Closely tracking the MUE could drive innovation to reduce the use of rare and critical metals and replace their use with less impactful materials. It could also be used to optimize the arrangement of components and their connectivity in a compute cabinet, as well as the sizing of the compute room and the arrangement of compute cabinets.

V. Limitations and Future Work

Due to limited disclosed manufacturer data, we used assumptions for volume estimates and may have missed elements containing copper or gold that could be similarly added.

This study focuses on the computing aspect of supercomputers, excluding the cooling infrastructure (cooling equipment and piping), the physical infrastructure (chassis, cabinets, metallic components, building), storage and management groups and their interconnections. The power supply components considered are limited to those within the compute room, excluding other infrastructures supplying power such as the generators and the ATS. Future work should include all these aspects for completeness.

The scope of this study is limited to copper and gold. Ideally, all resource-constrained material should be considered to build a comprehensive list of raw minerals required for supercomputer IT components manufacturing, feed an LCI, and have an accurate and exhaustive MUE metric.

We elaborate our models on high-end components used by HPC infrastructures. They may need to be adapted to less powerful IT components to be compatible with any data center.

Metal recycling is critical to ensure a sustainable development of such infrastructures. In future work, the MUE metric could include the proportion of metals that are recycled to better represent the abiotic resources depletion.

VI. Conclusion

Given the lack of data centers metrics accounting for the amount of metal required to manufacture these massive infrastructures, we introduce a novel metric to fill the gap, the Metal Usage Effectiveness, which we define as the ADP divided by the computing power of the data center or supercomputer.

Although progress is being made to advance the LCA of manufactured products and assess their environmental impact, the composition of high-end components is not publicly available. To alleviate this problem and to demonstrate that the MUE metric could be implemented in practice, we propose a set of models to estimate the amount of copper and gold contained in the main components of the Compute Room, in a multi-scale approach.

We apply our volume estimation models to the Frontier supercomputer to compute its MUE and provide a granular analysis of copper and gold location in its IT hardware components. Its component breakdown and their interconnections relies solely on publicly available information and images. The obtained weight proportions for RAM modules and PCB are in the right order of magnitude. Unlike related work that estimates metal composition at the recycling stage, our method allows this metric to be computed throughout the lifecycle of IT hardware components.

Currently, our estimation models and study of Frontier are limited to computing hardware and copper and gold. To compute a complete MUE, cooling, storage, and monitoring infrastructure should be included, as well as all minerals needed to produce a complete LCI. An accurate MUE value could be achieved in collaboration or provided by manufacturers.

VII. Acknowledgement

We thank Benoit Ruiz (CSR & Data Director, PICS Studio) and Rémi Bouzel (Scientific Director, Qarnot) for their advice and valuable insights. The authors acknowledge the financial support of Hi! PARIS Center (https://www.hi-paris.fr).

REFERENCES

[1] (2025) The multi-faceted challenge of powering ai. [Online]. Available: https://energy.mit.edu/news/the-multi-faceted-challenge-of-powering-ai/

[2] "Notice of final determination on 2023 DOE Critical Materials List," *Federal Register*, vol. 88 FR 51792, pp. 51 792–51 798, 2023.

[3] E. Commission, E. Directorate-General for Internal Market, Industry, SMEs, M. Grohol, and C. Veeh, *Study on the critical raw materials for the EU 2023 – Final report*. Publications Office of the European Union, 2023.

[4] (2024) Elementarium. [Online]. Available: https://lelementarium.fr/

[5] ADEME. (2024) Etude numérique et métaux. [Online]. Available: https://librairie.ademe.fr/economie-circulaire-et-dechets/7713-etude-numerique-et-metaux.html

[6] L. Wang and S. Khan, "Review of performance metrics for green data centers: A taxonomy study," *The Journal of Supercomputing*, vol. 63, pp. 1–18, 03 2011.

[7] M. Jamalzadeh and N. Behravan, "An exhaustive framework for better data centers' energy efficiency and greenness by using metrics," 2012. [Online]. Available: https://api.semanticscholar.org/CorpusID:110876212

[8] L. Siso *et al.*, "CoolEmAll D5.6 final metrics and benchmarks," IRIT-Institut de recherche en informatique de Toulouse, Research Report, Mar. 2014. [Online]. Available: https://hal.science/hal-01818070

[9] V. Reddy *et al.*, "Metrics for sustainable data centers," *IEEE Transactions on Sustainable Computing*, vol. 2, no. 3, pp. 290–303, 2017.

[10] C. Bekas and A. Curioni, "A new energy aware performance metric," *Computer Science - Research and Development*, vol. 25, pp. 187–195, 09 2010.

[11] M. Uddin, "Criteria to select energy efficiency metrics to measure performance of data centre," *International Journal of Energy Technology and Policy*, vol. 8, 06 2012.

[12] C. Fiandrino *et al.*, "Performance and energy efficiency metrics for communication systems of cloud computing data centers," *IEEE Transactions on Cloud Computing*, vol. 5, no. 4, pp. 738–750, 2017.

[13] X. Shao *et al.*, "A review of energy efficiency evaluation metrics for data centers," *Energy and Buildings*, vol. 271, p. 112308, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778822004790

[14] V. Avelar and D. Azevedo, "PUE™: A comprehensive examination of the metric," 2012. [Online]. Available: https://api.semanticscholar.org/CorpusID:167503232

[15] C. Belady and A. Rawson, "Green grid data center power efficiency metrics: PUE and DCIE," 2008. [Online]. Available: https://api.semanticscholar.org/CorpusID:270134891

[16] D. Azevedo, S. C. Belady, and J. Pouchet, "Water usage effectiveness (WUE): A green grid datacenter sustainability metric," *The Green Grid*, vol. 32, 2011.

[17] B. Subramaniam and W.-c. Feng, "The green index: A metric for evaluating system-wide energy efficiency in HPC systems," in *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*, May 2012, pp. 1007–1013.

[18] A. Aravanis *et al.*, "Metrics for assessing flexibility and sustainability of next generation data centers," in *2015 IEEE Globecom Workshops (GC Wkshps)*, 2015, pp. 1–6.

[19] G. Lykou *et al.*, "A new methodology toward effectively assessing data center sustainability," *Computers & Security*, vol. 76, pp. 327–340, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167404817302754

[20] E. Karanikolaou and M. Bekakos, "Action: A new metric for evaluating the energy efficiency on high performance computing platforms (ranked on Green500 List)," *WSEAS TRANSACTIONS ON COMPUTERS*, vol. 21, pp. 23–30, 01 2022.

[21] D. Azevedo, M. Patterson, J. Pouchet, and R. Tipley, "Carbon usage effectiveness (CUE): A green grid data center sustainability metric," *The green grid*, vol. 32, 2010.

[22] "New data center energy efficiency evaluation index," *DPPE Measurement Guidelines*, 2012.

[23] Emerson, "Recycling ratios: The next step for data center sustainability," 2011.

[24] M. Hoosain *et al.*, "Tools towards the sustainability and circularity of data centers," *Circ Econ Sustain*, vol. 3, pp. 173–197, 2023.

[25] J. Goddin *et al.*, "Circularity indicators: An approach to measuring circularity, methodology," 12 2019.

[26] R. U. Ayres, "Life cycle analysis: A critique," *Resources, Conservation and Recycling*, vol. 14, no. 3, pp. 199–223, 1995, life Cycle Management. [Online]. Available: https://www.sciencedirect.com/science/article/pii/092134499500017D

[27] J. Guinée *et al.*, "Life Cycle Assessment: Past, present, and future," *Environmental Science & Technology*, vol. 45, no. 1, pp. 90–96, 2011, pMID: 20812726. [Online]. Available: https://doi.org/10.1021/es101316v

[28] W. Klöpffer and B. Grahl, "Life cycle assessment (LCA): A guide to best practice," *Life Cycle Assessment (LCA): A Guide to Best Practice*, pp. 1–396, 04 2014.

[29] S. Hellweg and L. M. i Canals, "Emerging approaches, challenges and opportunities in life cycle assessment," *Science*, vol. 344, no. 6188, pp. 1109–1113, 2014. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.1248361

[30] J. B. Guinée and R. Heijungs, "A proposal for the definition of resource equivalency factors for use in product life-cycle assessment," *Environmental Toxicology and Chemistry*, vol. 14, no. 5, pp. 917–925, 05 1995. [Online]. Available: https://doi.org/10.1002/etc.5620140525

[31] R. Schulze *et al.*, "Abiotic resource use in life cycle impact assessment—part i- towards a common perspective," *Resources, Conservation and Recycling*, vol. 154, p. 104596, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921344919305026

[32] ——, "Abiotic resource use in life cycle impact assessment—part ii – linking perspectives and modelling concepts," *Resources, Conservation and Recycling*, vol. 155, p. 104595, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921344919305014

[33] (2024) Your source for LCA and sustainability data. [Online]. Available: https://nexus.openlca.org/databases

[34] Euromines. (2020) Key value chain electronics. [Online]. Available: https://euromines.org/files/key_value_chain_electronics_euromines_final.pdf/

[35] J. R. Centre, I. for Environment, Sustainability *et al.*, *Environmental footprint and material efficiency support for product policy – Report on benefits and impacts/costs of options for different potential material efficiency requirements for electronic displays*. Publications Office, 2013.

[36] C. Fitzpatrick *et al.*, "Conflict minerals in the compute sector: estimating extent of tin, tantalum, tungsten, and gold use in ict products." *Environmental science & technology*, vol. 49 2, pp. 974–81, 2015. [Online]. Available: https://api.semanticscholar.org/CorpusID:22162770

[37] J. Malmodin *et al.*, "A high-level estimate of the material footprints of the ICT and the E&M sector," in *ICT for Sustainability*, 05 2018.

[38] R. Charles *et al.*, "An investigation of trends in precious metal and copper content of ram modules in weee: Implications for long term recycling potential," *Waste Management*, vol. 60, pp. 505–520, 2017, special Thematic Issue: Urban Mining and Circular Economy. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0956053X16306778

[39] K. Ulman *et al.*, "Disentanglement of random access memory cards to regenerate copper foil: A novel thermo-electrical approach," *Waste Management*, vol. 81, pp. 138–147, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0956053X1830597X

[40] C. Gan *et al.*, "Technological sustainable materials and enabling in semiconductor memory industry: A review," *e-Prime - Advances in Electrical Engineering, Electronics and Energy*, vol. 5, p. 100245, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772671123001407

[41] S. Bobba *et al.*, "Critical raw materials for strategic technologies and sectors in the eu - a foresight study," 09 2020.

[42] J. H. Park and D. J. Fray, "Recovery of high purity precious metals from printed circuit boards oung," 2009. [Online]. Available: https://api.semanticscholar.org/CorpusID:621829

[43] W. Bizzo *et al.*, "Characterization of printed circuit boards for metal and energy recovery after milling and mechanical separation," *Materials*, vol. 73390, pp. 4555–4566, 06 2014.

[44] M. Arshadi *et al.*, "Content evaluation of different waste pcbs to enhance basic metals recycling," *Resources, Conservation and Recycling*, vol. 139, pp. 298–306, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921344918303070

[45] M. Bilesan *et al.*, "Efficient separation of precious metals from computer waste printed circuit boards by hydrocyclone and dilution-gravity methods," *Journal of Cleaner Production*, vol. 286, 12 2020.

[46] J. Allocco, "Laminated bus bars for power system interconnects," *Proceedings of APEC 97 - Applied Power Electronics Conference*, vol. 2, pp. 585–589 vol.2, 1997. [Online]. Available: https://api.semanticscholar.org/CorpusID:109985709

[47] (2023) What is 2D, 2.5D & 3D Packaging of integrated chips? [Online]. Available: https://techovedas.com/what-is-2d-2-5d-3d-packaging-of-integrated-chips/

[48] (2021) Scaling bump pitches in advanced packaging. [Online]. Available: https://semiengineering.com/scaling-bump-pitches-in-advanced-packaging/

[49] (2016) Copper pillar electroplating tutorial. [Online]. Available: https://www.dupont.com/blogs/copper-pillar-electroplating-tutorial.html

[50] (2025) What's Next For Through-Silicon Vias. [Online]. Available: https://semiengineering.com/whats-next-for-tsvs/

[51] (2021) How to Build Multilayer PCB Stack-up. [Online]. Available: https://www.protoexpress.com/blog/build-multilayer-pcb-stack-up/

[52] E. J. Wyrwas, "Body of knowledge for graphics processing units (gpus)," 2018. [Online]. Available: https://api.semanticscholar.org/CorpusID:69692511

[53] (2023) What is a pcb on a gpu? [Online]. Available: https://www.raypcb.com/gpu-pcb/

[54] (2014) Capacitor Dielectric Comparison. [Online]. Available: http://www.holystonecaps.com/PDF/TechNotes/20140116144107O.Capdielectriccomp3.pdf

[55] (2021) Chapter 22: Interconnects for 2d and 3d architectures. [Online]. Available: http://eps.ieee.org/hir

[56] (2024) Ssd form factors. [Online]. Available: https://www.snia.org/forums/cmsi/knowledge/formfactors/

[57] D. De Sensi *et al.*, "An in-depth analysis of the Slingshot interconnect," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, Nov. 2020, p. 1–14. [Online]. Available: http://dx.doi.org/10.1109/SC41405.2020.00039

[58] J. Kim *et al.*, "Technology-driven, highly-scalable dragonfly topology," *2008 International Symposium on Computer Architecture*, pp. 77–88, 2008. [Online]. Available: https://api.semanticscholar.org/CorpusID:79473

[59] K. Ahmed *et al.*, "A review of data centers energy consumption and reliability modeling," *IEEE Access*, vol. PP, pp. 1–1, 11 2021.

[60] (2024) Datacenter Anatomy Part 1: Electrical Systems. [Online]. Available: https://semianalysis.com/2024/10/14/datacenter-anatomy-part-1-electrical/

[61] J. Dongarra and A. Geist, "Report on the oak ridge national laboratory's Frontier system," *Univ. of Tennessee, Knoxville, Tech. Rep. Tech Report No. ICL-UT-22-05*, 2022.

[62] S. Atchley *et al.*, "Frontier: Exploring exascale," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3581784.3607089

[63] W. Brewer *et al.*, "A digital twin framework for liquid-cooled supercomputers as demonstrated at exascale," in *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2024, pp. 1–18.

[64] W. Doll *et al.*, "Flexible and adaptable computing system infrastructure," U.S. Patent US16/499,209, Mar. 27, 2018. [Online]. Available: https://patents.google.com/patent/US10842040B2/

[65] (2024) HPE Cray EX Supercomputer QuickSpecs. [Online]. Available: https://www.hpe.com/psnow/doc/a00094635enw

[66] (2023) Socket sp3 - packages - amd. [Online]. Available: https://en.wikichip.org/wiki/amd/packages/socket_sp3/

[67] (2021) Amd radeon instinct mi250x. [Online]. Available: https://www.techpowerup.com/gpu-specs/radeon-instinct-mi250x.c3837

[68] (2015) Ddr4 sdram rdimm. [Online]. Available: https://sg-repo-production-photos.s3.eu-central-1.amazonaws.com/aikido/cache/3b68b7e98bc36c6284d027ad8913b751.pdf

[69] (2024) HPE Cray EX Hardware Management Administration Guide (1.4). [Online]. Available: https://support.hpe.com/hpesc/public/docDisplay?docId=a00115093en_us\&page=Shasta_Hardware_Architecture.html

[70] (2024) HPE Slingshot Hardware Guide (2.1.3) (S-9001). [Online]. Available: https://support.hpe.com/hpesc/public/docDisplay?docId=dp00004988en_us\&docLocale=en_US

[71] (2022) Oam. [Online]. Available: https://www.molex.com/content/dam/molex/molex-dot-com/products/automated/en-us/salesdrawingpdf/209/209311/2093111115_sd.pdf?inline

[72] (2018) QSFP-DD hardware specification. [Online]. Available: http://www.qsfp-dd.com/wp-content/uploads/2018/09/QSFP-DD-Hardware-rev4p0-9-12-18-clean

[73] (2016) Examax. [Online]. Available: https://cdn.amphenol-cs.com/media/wysiwyg/files/drawing/10121067.pdf

[74] (2016) Gen 4 pcie connector & channel design and optimization: 16gt/s for free. [Online]. Available: https://www.oldfriend.url.tw/article/IEEE_paper/SI/Intel_PCIe%20connector%2016G_DesignCon2016_p89~.pdf

[75] (2019) Pci express flash drive. [Online]. Available: https://www.mouser.ch/datasheet/2/24/AS228AP2_Spec_Sheet-3475234.pdf

[76] (2023) Micron 7450 ssd series technical product specification. [Online]. Available: https://www.micron.com/content/dam/micron/global/public/documents/products/technical-marketing-brief/7450-nvme-ssd-tech-prod-spec.pdf

[77] D. Roweth *et al.*, "Hpe slingshot launched into network space."

[78] (2020) Inside rosetta: The engine behind cray's slingshot exascale-era interconnect. [Online]. Available: https://fuse.wikichip.org/news/3293/inside-rosetta-the-engine-behind-crays-slingshot-exascale-era-interconnect/

[79] (2024) Global general server pcb market research report 2024. [Online]. Available: https://www.marketresearchreports.com/gir/global-server-pcb-market-2024-manufacturers-regions-type-and-application-forecast-2030

[80] Thin build up substrate. [Online]. Available: https://www.lincstech.com/english/product/buildup/

[81] (2021) Integrating and operating hbm2e memory. [Online]. Available: https://www.micron.com/content/dam/micron/global/public/products/technical-marketing-brief/micron-hbm2e-memory-wp.pdf

[82] (2018) Ddr4 sdram memory. [Online]. Available: https://download.semiconductor.samsung.com/resources/product-guide/DDR4_Product_guide_May.18.pdf

[83] (2024) Stackup for 4, 6, 8,..., 18 layers multi-layer laminated structure. [Online]. Available: https://www.pcbway.com/multi-layer-laminated-structure.html

[84] S. Enoch *et al.*, "Design considerations for a new generation of sipms with unprecedented timing resolution," *Journal of Instrumentation*, vol. 16, pp. P02 019–P02 019, 02 2021.

[85] (2022) Amd, intel display chiplet packaging prowess. [Online]. Available: https://www.semiconductor-digest.com/amd-intel-display-chiplet-packaging-prowess/

[86] (2021) Enabling heterogenous integration through chiplet architectures. [Online]. Available: https://imapsource.org/api/v1/articles/128222-enabling-heterogenous-integration-through-chiplet-architectures.pdf

[87] T. M. Braun *et al.*, "Simulation of copper electrodeposition in millimeter size through-silicon vias," *Journal of The Electrochemical Society*, vol. 167, no. 16, p. 162508, dec 2020. [Online]. Available: https://dx.doi.org/10.1149/1945-7111/abd112

[88] (2024) PCB layers: Everything you need to know. [Online]. Available: https://www.wevolver.com/article/pcb-layers

[89] (2019) BCM56980 12.8 tb/s multilayer switch. [Online]. Available: https://docs.broadcom.com/doc/56980-DS

[90] (2020) Ddr4 point-to-point design guide. [Online]. Available: https://www.mouser.com/pdfDocs/Micron_DDR4_Design_Guide.pdf

[91] Leiton. Pcb weight calculation. [Online]. Available: https://www.leiton.de/leiton-tools-weight-calculation.html

[92] (2023) It takes 6000 gallons of water to cool the world's fastest supercomputer. [Online]. Available: https:https://en.wikichip.org/wiki/amd/packages/socket_sp3tionsOffice