



(12) APPLICATION

(11) 20230136

(13) A1

NORWAY

(19) NO

(51) Int Cl.

H05K 7/20 (2006.01)

G06F 9/50 (2006.01)

G06F 1/20 (2006.01)

G06F 1/32 (2006.01)

G06F 9/455 (2006.01)

Norwegian Industrial Property Office

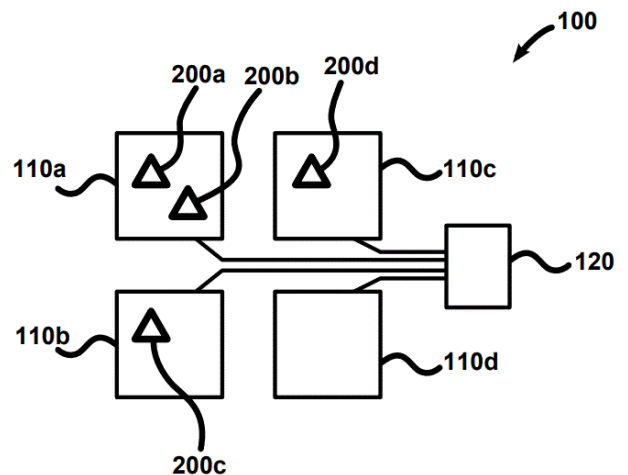
(21)	Application nr	20230136	(86)	Int. application day and application nr
(22)	Application day	2023.02.10	(85)	Entry into national phase
(24)	Date from which the industrial right has effect	2023.02.10	(30)	Priority
(41)	Available to the public	2024.03.25		
(71)	Applicant	Figurati AS, Gosenstien 17b, 4041 HAFRSFJORD, Norge		
(72)	Inventor	Richard Rettedal, Porsahagen 6, 4316 SANDNES, Norge Bjørn Knudsen, Gosenstien 17B, 4041 HAFRSFJORD, Norge Christopher Birkeland, Tastatunet 8, 4027 STAVANGER, Norge Matt Chauncy-Lie, Muségata 73, 4010 STAVANGER, Norge		
(74)	Agent or Attorney	HÅMSØ PATENTBYRÅ AS, Postboks 9, 4068 STAVANGER, Norge		

(54) Title **SYSTEM FOR AND METHOD OF HARNESSING HEAT GENERATED FROM RUNNING AT LEAST ONE VIRTUAL OPERATIVE SYSTEM INSTANCE**

(57) Abstract

A system and method for harnessing heat generated from running virtual operative system instances. The system includes networked host computers for running the one virtual operative system instances, a subsystem for harnessing heat generated by the host computers and one or more of the host computers configured to operate as a controller. The controller is configured with the steps of:

- for each instance, estimating a heat rate that will result from running the instance as a guest of a host computer;
- calculating an arrangement of the instances over the host computers so that a heat rate achieved by the host computers is maximized; and
- controlling the host computers to run the instances as defined in the calculated arrangement.



SYSTEM FOR AND METHOD OF HARNESSING HEAT GENERATED FROM RUNNING AT LEAST ONE VIRTUAL OPERATIVE SYSTEM INSTANCE

The present invention relates to a system for harnessing heat generated from running at least one virtual operative system instance. The present invention also relates to a
5 method of harnessing heat generated from running at least one virtual operative system instance.

Background

Nowadays, sustainability is becoming an increasing priority in the data center industry. For example, some known developments aim to harness heat (aka. to harvest heat)
10 generated by running computers in a data center and then have the harnessed heat be supplied to an external facility. The harnessed heat can be used for heating buildings (e.g. an apartment building, a public swimming pool facility) or food production processes (e.g. greenhouses) that require heat.

There are various known solutions for harnessing heat from computers used in a data
15 center. For example, a known heat harnessing system is based on a direct contact evaporative cooling system (e.g. ZutaCore HyperCool²[™]), in which at least one heat generating chip (e.g. a Central Processing Unit chip) is equipped with a device for transferring heat from the chip to a coolant (e.g. 3M[™] Novec[™] 7000 Engineered Fluid). The heat ensuing from the heat transfer is then transmitted to and used at the external
20 facility. Other heat harnessing systems are also known.

In this context, it observed that it can be challenging to orchestrate instances of virtual operative systems as guests of at least one host computer, when the at least one host computer is being used as a source of heat for a heat harnessing system.

Running instances of virtual operative systems as guests of a host computer allows the hardware resources of the host computer to be shared with the instances. Known approaches for running an instance of a virtual operative system include: a hypervisor software (e.g. KVM), which runs at least one virtual machine instance on the host computer; and a container engine software (e.g. Docker), which runs at least one container instance on the host computer. Due to each instance requiring computational resources in a way that varies over time, the overall requirements for hardware resources of the host computers change over time.

In know art, "orchestration" refers to methods of monitoring and managing in which of the at least one host computer is each instance of a virtual operative system being run. Several factors may be taken into account when orchestrating (virtual machine or container) instances. Typically, known orchestration methods aim to cool down hardware as efficiently as possible while achieving an even and low computing load on all host computers. Known orchestration methods may aim to even out server load in order to avoid the risk of overheating the hardware, which could lead to damage and/or failure. The known orchestration methods typically achieve low rates of the harnessed heat relative to the consumption of hardware resources by the instances in the at least one host computer.

Summary

The invention will now be disclosed and has for its object to remedy or to reduce at least one of the drawbacks of the prior art, or at least provide a useful alternative to prior art. The object is achieved through features, which are specified in the description below and in the claims that follow. The invention is defined by the independent claims. The dependent claims define advantageous embodiments of the invention.

According to a first aspect of the invention, there is provided a system for harnessing heat generated from running at least one virtual operative system instance. The system comprises:

- at least one host computer for running the at least one virtual operative system instance, the at least one host computer being networked; and

- a subsystem for harnessing heat generated by the at least one host computer.

One or more of the at least one host computer is configured to operate as a controller of the at least one host computer. The controller is configured to carry out the steps of:

- for each virtual operative system instance, estimating a heat rate that will result from running the instance as a guest of a host computer;

- calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that a heat rate achieved by the at least one host computer is maximized; and

- controlling the at least one host computer to run the at least one virtual operative system instance as defined in the calculated arrangement.

Thus, the system achieves a maximization of the heat rate generated by the at least one host computer in use relative to the number of virtual operative system instances being run. Thus, the residual heat levels per host computer are increased, which makes the at least one host computer more useful for the purpose of heat harnessing. Also, the system is advantageous in that the amount of host machines needed for the same workloads (caused by the virtual operative system instances) is reduced.

Optionally, for each host computer, the controller is configured with a target minimum heat rate to be achieved by the host computer, and the step of calculating an arrangement comprises the step of calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that an amount of host computers achieving the respective target minimum heat rate is maximized. Thus, the system can be configured to aim for a target minimum heat rate to be achieved by a host computer when the host computer is controlled to run a virtual operative system instance. This embodiment is advantageous, for example, when the subsystem for harnessing heat generated by the at least one host computer has an improved harnessing efficiency that depends on the host computer generating heat above a defined amount (e.g. an amount of heat sufficient to cause the evaporation of a coolant fluid). Thus, when calculating an arrangement of the at least one virtual operative system instance over the at least one host computer, the system can optimize the calculation based on the heat harnessing capabilities of the subsystem. Therefore, significant amounts of heat may be

harnessed over time, the harnessed heat being usable in an external facility.

Optionally, the step of estimating a heat rate comprises estimating the heat rate based on any of:

- 5 - previous data about a heat rate that resulted from running the virtual operative system instance as a guest of a host computer;
- real-time data about a heat rate that is resulting from running the virtual operative system instance as a guest of a host computer;
- data generated by a mathematical model of a heat rate that results from running the virtual operative system instance as a guest of a host computer; and/or
- 10 - data generated by a computer implemented simulation of a heat rate that results from running the virtual operative system instance as a guest of a host computer.

Optionally, the step of controlling the at least one host computer to run the at least one virtual operative system instance comprises the step of migrating a virtual operative system instance from one host computer to a different host computer.

- 15 Optionally, the controller is further configured to carry out any of the steps of:
 - shutting down a host computer that has not been controlled to run at least one virtual operative system instance; or
 - controlling a host computer that has not been controlled to run at least one virtual operative system instance so that the host computer runs in a mode for saving energy.
- 20 Thus, the number of resources running idle per host computer is minimized.

Optionally, the at least one virtual operative system instance comprises any of: a virtual machine instance; and/or a container instance.

- According to a second aspect of the invention, there is provided a method of harnessing heat generated from running at least one virtual operative system instance. The method
- 25 comprises the steps of:

- providing at least one host computer for running the at least one virtual operative system instance, the at least one host computer being networked;
- providing a subsystem for harnessing heat generated by the at least one host computer;

- configuring one or more of the at least one host computer to operate as a controller of the at least one host computer;
 - for each virtual operative system instance, estimating a heat rate that will result from running the instance as a guest of a host computer;
 - 5 - calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that a heat rate achieved by the at least one host computer is maximized; and
 - controlling the at least one host computer to run the at least one virtual operative system instance as defined in the calculated arrangement.
- 10 Optionally, the method comprises the step of:
- for each host computer, configuring the controller with a target minimum heat rate to be achieved by the host computer, and
- wherein the step of calculating an arrangement comprises the step of:
- calculating an arrangement of the at least one virtual operative system instance over the
 - 15 at least one host computer so that an amount of host computers achieving the respective target minimum heat rate is maximized.

Optionally, the step of estimating a heat rate comprises estimating the heat rate based on any of:

- previous data about a heat rate that resulted from running the virtual operative system
- 20 instance as a guest of the host computer;
- real-time data about a heat rate that is resulting from running the virtual operative system instance as a guest of a host computer;
- data generated by a mathematical model of a heat rate that results from running the virtual operative system instance as a guest of a host computer; and/or
- 25 - data generated by a computer implemented simulation of a heat rate that results from running the virtual operative system instance as a guest of a host computer.

Optionally, the step of controlling the at least one host computer to run the at least one virtual operative system instance comprises the step of migrating a virtual operative system instance from one host computer to a different host computer.

Optionally, the controller is further configured to carry out any of the steps of:

- shutting down a host computer that has not been controlled to run at least one virtual operative system instance; or
- controlling a host computer that has not been controlled to run at least one virtual operative system instance so that the host computer runs in a mode for saving energy.

Optionally, the at least one virtual operative system instance comprises any of: a virtual machine instance; and/or a container instance.

Brief description of the drawings

In the drawings:

- 10 Fig. 1 shows schematic view of a system embodiment including four host computers, the four host computers running four virtual operative system instances;
- Fig. 2 shows a schematic view of a one graph illustrating the distribution of computational load over four host computers being controlled to run virtual operative system instances as defined in one arrangement;
- 15 Fig. 3 shows a schematic view of another graph illustrating the distribution of computational load over four host computers being controlled to run virtual operative system instances as defined in another arrangement.

Detailed description

- 20 In the figures, same or corresponding elements are indicated by same reference numerals. For clarity reasons, some elements may in some of the figures be without reference numerals. A person skilled in the art will understand that the figures are just principal drawings. The relative proportions of individual elements may also be distorted.

- Turning now to Fig. 1, it shows a system embodiment 100 including four host computers 110a–110d for running four virtual operative system instances 200a–200d. The system 100 also includes a subsystem 120 for harnessing heat generated by the four host computers 110a–110d. It will be understood by the skilled person that the system 100

shown in Fig. 1 is illustrated in a simplified manner for the purpose of illustrating an embodiment of the invention. In other embodiments, the number, structure, form and/or organization of the at least one host computer and the heat harnessing subsystem may vary substantially and within what is known in the data center practice.

- 5 The host computers 110a–110d may be embodied in a manner known from the data center practice. For example, a data center may be provided in which a number of host computers is installed in various racks and with connection to known systems available in a typical data center, such as a power system, a network, an air-cooling system, among others. The host computers may be configured to operate software suitable for running
10 virtual operative system instances, such as a hypervisor software (e.g. KVM) or a container engine software (e.g. Kubernetes).

The subsystem 120 for harnessing heat from the host computers 110a-110d may be provided in many known manners. For example, a known subsystem is based on a direct contact evaporative cooling system using a two-phase coolant (e.g. 3M™ Novec™ 7000
15 Engineered Fluid), the coolant having a pre-defined boiling temperature. The harnessing of heat is achieved by transmitting the coolant, in liquid phase, onto direct contact with one or more chips of a host computer 110a-110d. When that contact happens and the chip generates sufficient heat to cause the coolant to boil, the coolant evaporates and flows, in gaseous form, towards the subsystem 120. At that stage, the gaseous, hot
20 coolant is used as a source of heat to be harnessed. The heat is harnessed and transmitted to an external facility for reuse. The harnessing of the heat from the coolant lowers the temperature of the coolant, which in turn causes the gaseous coolant to change back to the liquid phase. Harnessing heat from the hot, gaseous coolant may be achieved in many known ways, for example by a heat exchanger device in which the heat
25 from the coolant is transferred to a separate circuit including a fluid (e.g. water) for transmitting the harnessed heat to an external facility. It will be understood that other subsystems 120 for harnessing heat are known and may be used for harnessing heat from one or more host computers.

In the system 100, the host computer 110a shown at the top left corner of Fig. 1, is

additionally configured to operate as a controller for controlling the running of the four virtual operative system instances 200a-200d over the four host computers 110a-110d. It will be understood that the controller may be implemented in other known ways, including having the controller be executed in a cooperative or decentralized manner by two or more host computers, among other known approaches. It will also be understood that a host computer 110a-110d is configurable to both operate as the controller and run one or more virtual operative system instances.

In the embodiment shown in Fig. 1, the controller has controlled the four host computers 110a-110d to run the four virtual operative system instances 200a-200d as shown: the host computer 110a on the top-left corner is controlled to run two of the instances 200a,200b; the host computer 110b on the lower-left corner is controlled to run one instance 200c; and the host computer 110c on the top-right corner is controlled to run one instance 200d. This arrangement of the four virtual operative system instances 200a-200d over the four host computers 110a-110d was calculated so that a heat rate achieved by the four host computers 110a-110d is maximized. When calculating the arrangement, the controller carried out the step of estimating, for each virtual operative system instance 200a-200d, a heat rate that will result from running the instance as a guest of a host computer 110a-110d. This estimation of a heat rate may be implemented in many known ways. For example, the estimation may involve estimating the percentage of a period of time during which a temperature of a chip (e.g. CPU) in a host computer will be above a certain minimum temperature. Or, for example, the estimation may be implemented by estimating a heat rate in kWh. It will be understood that other known ways of estimating a heat rate may be used.

In some embodiments, the heat rate estimation is based on any of:

- previous data about a heat rate that resulted from running the virtual operative system instance as a guest of a host computer;
- real-time data about a heat rate that is resulting from running the virtual operative system instance as a guest of a host computer;
- data generated by a mathematical model of a heat rate that results from running the virtual operative system instance as a guest of a host computer; and/or

- data generated by a computer implemented simulation of a heat rate that results from running the virtual operative system instance as a guest of a host computer.

These data options may be chosen exclusively or non-exclusively combined in many ways. Moreover, these data options may be combined with other known data options.

5 After the estimating step, the controller carries out a step of calculating the arrangement. The skilled person will know many search algorithms for efficiently finding an arrangement in which the heat rate achieved by the four host computers 110a-110d is maximized. Also, it will be appreciated that the skilled person will know how a plurality of heat rate estimates may be accumulated to represent a plurality of respective instances
10 running as guests of the same host computer.

For example, in Fig. 1, the host computers 110b,110c on the lower-left corner and the top-right corner have been controlled to run only one virtual operative system instance 200c,200d each. This portion of the arrangement resulted from estimating that each of the instances 200c,200d will cause a host computer to generate a heat rate sufficiently
15 high to justify having only one instance running in each of the two host computers 110b,110c.

Fig. 2 shows, for a system embodiment, a graph illustrating the distribution of computational loads 110a'-110d' over four host computers being controlled to run virtual operative system instances as calculated in one arrangement. The two host computers
20 related to the computational loads 110a',110b' shown on the left-hand side of Fig. 2 have been assigned with all the instances to be run, whereas the two host computers related to the computational loads 110c',110d' shown on the right-hand side of Fig. 2 have not been assigned to run any instances. Also, the later host computers, i.e. related to the computational loads 110c',110d' on the right-hand side, have been shut down in order to
25 reduce the power consumption. Thus, the rate of generated heat by the four host computers is maximized.

Fig. 3 shows, for a system embodiment, another graph illustrating the distribution of computational loads 110a''-110d'' over four host computers being controlled to run virtual operative system instances as calculated in another arrangement.

In Fig. 3, the controller of the system embodiment is, for each of the four host computers, configured with a target minimum heat rate to be achieved by the host computer. These four target minimum heat rate configurations are shown in Fig. 3 in the form of a dashed horizontal line intersecting the computational loads 110a''-110d'', thus representing that, in this particular embodiment, the controller is configured with the same target minimum heat rate for each of the four host computers. It will be understood that different target minimum heat rates may be configured for each host computer, and this configuration may take into account various parameters related to the heat generating capability of the host computer.

When calculating an arrangement of the at least one virtual operative system instance over the four host computers, the controller is configured to maximize an amount of host computers achieving the respective target minimum heat rate. As shown in Fig. 3, all the host computers running instances are achieving computational loads 110a''-110c'' above the target minimum heat rates, even though a few of them may still have capacity to run further instances.

It should be noted that the above-mentioned embodiments illustrate rather than limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. Use of the verb "comprise" and its conjugations does not exclude the presence of elements or steps other than those stated in a claim. The article "a" or "an" preceding an element does not exclude the presence of a plurality of such elements.

The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage.

The invention may be implemented by means of hardware comprising several distinct elements, and by means of a suitably programmed computer. In the system claim listing several means, several of these means may be embodied by one and the same item of hardware.

C l a i m s

1. A system for harnessing heat generated from running at least one virtual operative system instance, the system comprising:
 - at least one host computer for running the at least one virtual operative system instance, the at least one host computer being networked; and
 - a subsystem for harnessing heat generated by the at least one host computer, wherein one or more of the at least one host computer is configured to operate as a controller of the at least one host computer, and wherein the controller is configured to carry out the steps of:
 - for each virtual operative system instance, estimating a heat rate that will result from running the instance as a guest of a host computer;
 - calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that a heat rate achieved by the at least one host computer is maximized; and
 - controlling the at least one host computer to run the at least one virtual operative system instance as defined in the calculated arrangement.
2. System according to claim 1, wherein, for each host computer, the controller is configured with a target minimum heat rate to be achieved by the host computer, and wherein the step of calculating an arrangement comprises the step of:
 - calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that an amount of host computers achieving the respective target minimum heat rate is maximized.
3. System according to any of the preceding claims, wherein the step of estimating a heat rate comprises estimating the heat rate based on any of:
 - previous data about a heat rate that resulted from running the virtual operative system instance as a guest of a host computer;
 - real-time data about a heat rate that is resulting from running the virtual operative system instance as a guest of a host computer;

- data generated by a mathematical model of a heat rate that results from running the virtual operative system instance as a guest of a host computer; and/or
- data generated by a computer implemented simulation of a heat rate that results from running the virtual operative system instance as a guest of a host computer.

- 5 4. System according to any of the preceding claims, wherein the step of controlling the at least one host computer to run the at least one virtual operative system instance comprises the step of migrating a virtual operative system instance from one host computer to a different host computer.
- 10 5. System according to any of the preceding claims, wherein the controller is further configured to carry out any of the steps of:
- shutting down a host computer that has not been controlled to run at least one virtual operative system instance; or
 - controlling a host computer that has not been controlled to run at least one virtual operative system instance so that the host computer runs in a mode for
- 15 saving energy.
6. System according to any of the preceding claims, wherein the at least one virtual operative system instance comprises any of: a virtual machine instance; and/or a container instance.
- 20 7. A method of harnessing heat generated from running at least one virtual operative system instance, the method comprising the steps of:
- providing at least one host computer for running the at least one virtual operative system instance, the at least one host computer being networked;
 - providing a subsystem for harnessing heat generated by the at least one host
- 25 computer;
- configuring one or more of the at least one host computer to operate as a controller of the at least one host computer;
 - for each virtual operative system instance, estimating a heat rate that will result from running the instance as a guest of a host computer;

- calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that a heat rate achieved by the at least one host computer is maximized; and

- controlling the at least one host computer to run the at least one virtual operative system instance as defined in the calculated arrangement.

8. Method according to claim 7, the method comprises the step of:

- for each host computer, configuring the controller with a target minimum heat rate to be achieved by the host computer, and

wherein the step of calculating an arrangement comprises the step of:

- calculating an arrangement of the at least one virtual operative system instance over the at least one host computer so that an amount of host computers achieving the respective target minimum heat rate is maximized.

9. Method according to any of the claims 7 to 8, wherein the step of estimating a heat rate comprises estimating the heat rate based on any of:

- previous data about a heat rate that resulted from running the virtual operative system instance as a guest of the host computer;

- real-time data about a heat rate that is resulting from running the virtual operative system instance as a guest of a host computer;

- data generated by a mathematical model of a heat rate that results from running the virtual operative system instance as a guest of a host computer; and/or

- data generated by a computer implemented simulation of a heat rate that results from running the virtual operative system instance as a guest of a host computer.

10. Method according to any of the claims 7 to 9, wherein the step of controlling the at least one host computer to run the at least one virtual operative system

instance comprises the step of migrating a virtual operative system instance from one host computer to a different host computer.

11. Method according to any of the claims 7 to 10, wherein the controller is further configured to carry out any of the steps of:

- shutting down a host computer that has not been controlled to run at least one

virtual operative system instance; or

- controlling a host computer that has not been controlled to run at least one virtual operative system instance so that the host computer runs in a mode for saving energy.

- 5 12. Method according to any of the claims 7 to 11, wherein the at least one virtual operative system instance comprises any of: a virtual machine instance; and/or a container instance.

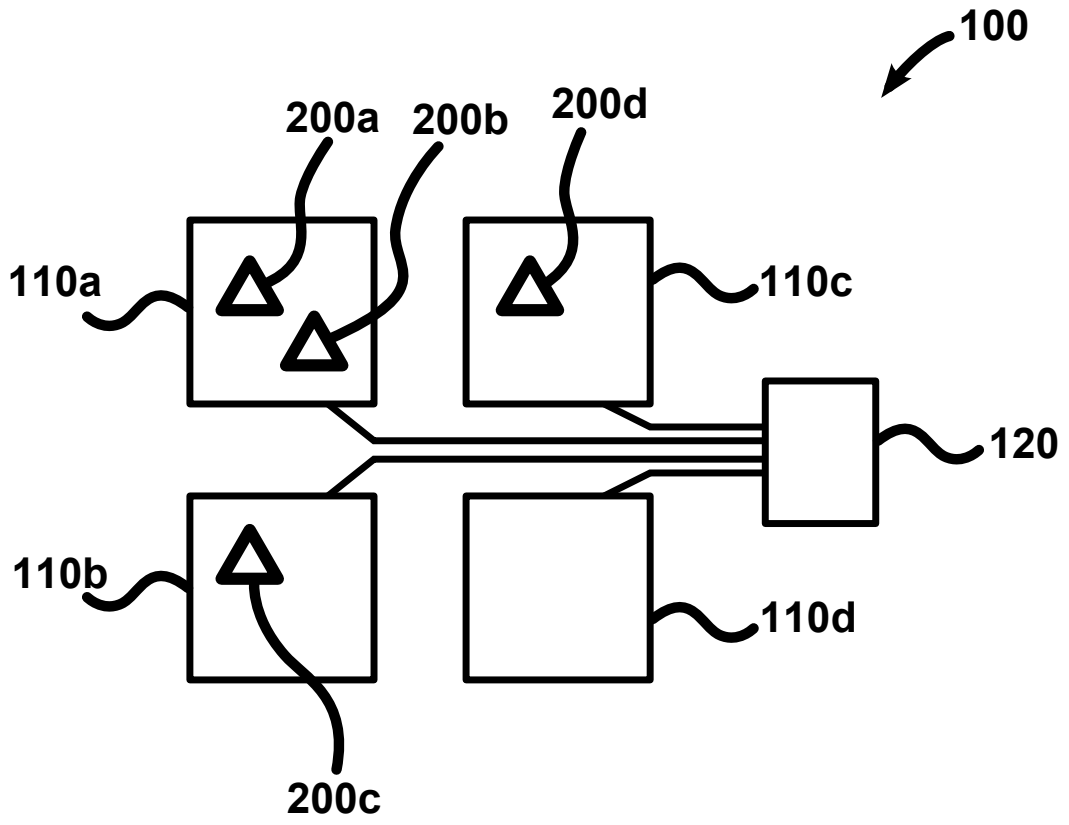


Fig. 1

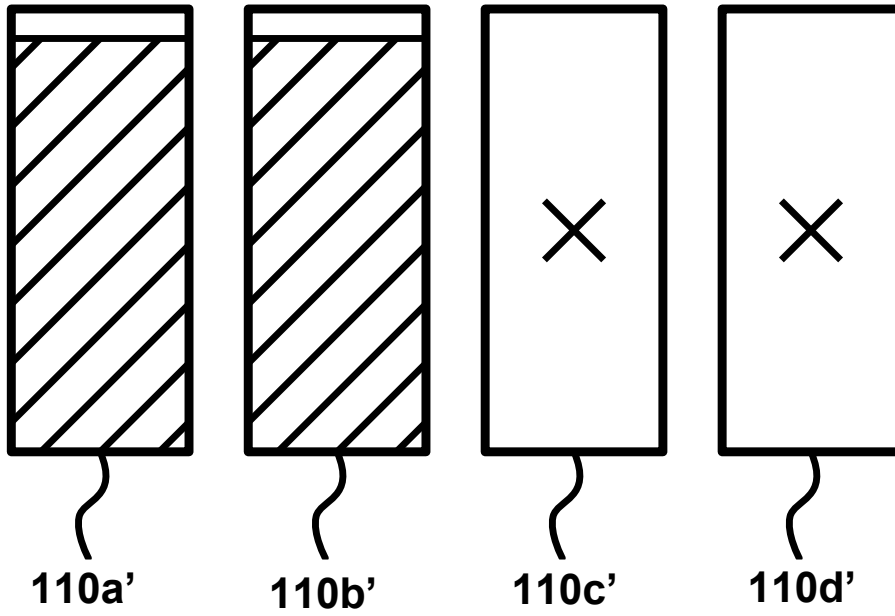


Fig. 2

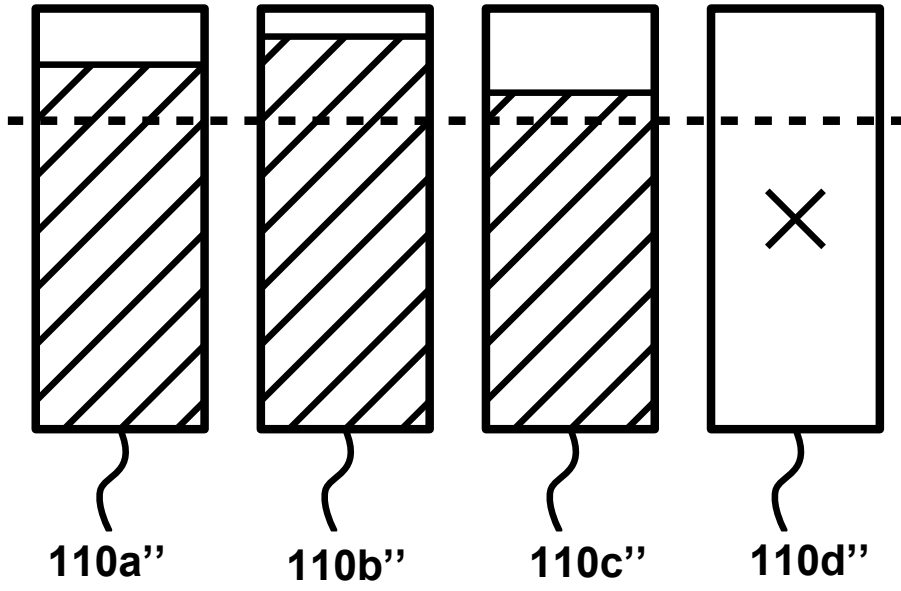


Fig. 3