



- (51) International Patent Classification:
G06F 11/07 (2006.01)
- (21) International Application Number:
PCT/US2016/019765
- (22) International Filing Date:
26 February 2016 (26.02.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
201510088768.3 26 February 2015 (26.02.2015) CN
- (71) Applicant: **ALIBABA GROUP HOLDING LIMITED**
[—/US]; Fourth Floor, One Capital Place, P.O. Box 847,
George Town, Grand Cayman (KY).
- (72) Inventor: **FEI, Hui**; Alibaba Group Legal Department, 5/f,
Building 3, No. 969, West Wen Yi Road, Yu Hang Dis-
trict, Hangzhou 311121 (CN).
- (74) Agent: **MURABITO, Anthony, C.**; Murabito, Hao &
Barnes, Llp, 2 North Market Street, 3rd Floor, San Jose,
CA 95113 (US).
- (81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every
kind of regional protection available*): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR PREDICTING GPU MALFUNCTIONS




FIG. 1

(57) Abstract: A method of predicting GPU malfunctions includes installing a daemon program at a GPU node, the daemon program periodically collecting GPU status parameters corresponding to the GPU node at a pre-determined time period. The method also includes obtaining the GPU status parameters from the GPU node and comparing the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, where the mean status fault parameters are obtained by use of a pre-configured statistical model. Prior to a GPU enters a malfunction state, the GPU can be replaced, or the programs executing on the GPU can be migrated to other GPUs for execution, without affecting the normal business operations.



LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, **Published:**

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *with international search report (Art. 21(3))*
GW, KM, ML, MR, NE, SN, TD, TG).

METHOD AND APPARATUS FOR PREDICTING GPU MALFUNCTIONS

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefits to Chinese Patent Application No. 201510088768.3, filed on Feb. 26, 2015, which is incorporated herein by reference in its entirety.

TECHNICAL FIELD

[0002] The present invention relates generally to the field of computer communication technologies, and more particularly to a method and apparatus for predicting GPU malfunctions.

BACKGROUND

[0003] With the development of the computer communication technologies, the utilization of graphics processing units (GPUs) in the general computing field is becoming wider and wider. Equipped with hundreds of computing cores, a GPU can achieve a computing power in terms of tera floating-points operations per second (TFLOPS). For general computing purposes, the powerful floating-points computation capacity of a GPU is beyond comparable with that of a CPU. Therefore, the general computing power of a GPU can be utilized to compensate the deficiency of a CPU in parallel computing.

[0004] In order to monitor the status of each GPU of a GPU cluster, the current technology deploys daemon programs for each GPU node, the daemon program collecting the GPU information about such as the GPU model, temperature, power consumption, usage duration, usage status, etc. The daemon program also displays the collected GPU information. Based on the collected GPU information, it is determined whether there are occurrences of errors or failures with regard to the GPU; and if so, alerts are generated accordingly.

[0005] With the current technology, it is only upon the detection of faults with a GPU that a user is alerted with the GPU malfunctioning. It is also at this point of time that the user replaces the faulty GPU, or migrates the programs executing on the faulty GPU to some other normal functioning GPUs for execution, imposing negative impact on the normal business operations.

SUMMARY

[0006] In order to solve the current technical problem, the present disclosure provides for a method and apparatus for predicting GPU malfunctions. Prior to a GPU enters a malfunction state, status parameters and mean status fault parameters are utilized to determine whether a GPU is about to malfunction such that a malfunction state of the GPU can be predicted. Consequently, prior to the GPU enters a malfunction state, the GPU can be replaced, or the programs executing on the GPU can be migrated to other GPUs for execution, without affecting the normal business operations.

[0007] According to an exemplary embodiment of the present disclosure, a method of predicting GPU malfunctions includes installing a daemon program at a GPU node, the daemon program periodically collecting the GPU status parameters corresponding to the GPU node at a pre-determined time period. The method also includes obtaining the GPU status parameters from the GPU node. The method further includes comparing the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, where the mean status fault parameters are obtained by use of pre-configured statistical models.

[0008] Alternatively, when the status parameter is temperature, the comparing of the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction includes performing statistics to obtain a temperature count of a number of times the GPU incurs a temperature greater than a pre-determined temperature threshold and a GPU temperature standard deviation threshold for the GPU. The comparing also includes comparing the obtained temperature count with a mean temperature fault count and comparing a GPU temperature fault standard deviation with the GPU temperature standard deviation threshold. The comparing further determines that, if the temperature count is greater than the mean temperature fault count, and if the GPU temperature fault standard deviation is less than the GPU temperature standard deviation threshold, the GPU is to malfunction. Otherwise, if the temperature count is less than the mean temperature fault count, or if the GPU temperature fault standard deviation is greater than the GPU temperature standard deviation threshold, the comparing determines that the GPU is not to malfunction.

[0009] Further alternatively, the GPU status parameters collected from the GPU node include a GPU model, temperature, and usage status. The mean status fault parameters are obtained by use of a pre-configured statistical model through the following steps. First, it is determined, based on the usage status, whether the GPU malfunctions. Second, in response to a determination that the GPU does not malfunction, based upon the GPU model, the temperature collected from the GPU node is

stored in an information storage space corresponding to the GPU model. In response to a determination that the GPU does malfunction, based on the GPU model, the stored temperatures are obtained from an information storage space corresponding to the GPU model. Based on the temperature collected from the GPU node and the stored temperatures obtained from the storage space corresponding to the GPU model, statistics are performed to compute an arithmetic mean temperature fault count and a GPU temperature fault standard deviation, by use of a pre-configured temperature statistical model.

[0010] Still further alternatively, the pre-configured temperature statistical model includes a mean temperature fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU obtained from the information storage space corresponding to the GPU model. The pre-configured temperature statistical model also includes a temperature fault standard deviation model configured to compute a temperature fault standard deviation based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, and the mean temperature fault count.

[0011] Yet further alternatively, when the status parameter is power consumption, the comparing of the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction includes performing statistics to obtain a power consumption count of a number of times the GPU incurs power consumption greater than a pre-determined power consumption threshold and a GPU power consumption standard deviation threshold for the GPU. The comparing also includes comparing the obtained power consumption count with a mean power consumption fault count and comparing a GPU power consumption fault standard deviation with the GPU power consumption standard deviation threshold. The comparing further determines that, if the power consumption count is greater than the mean power consumption fault count, and if the GPU power consumption fault standard deviation is less than the GPU temperature standard deviation threshold, the GPU is to malfunction. Otherwise, if the temperature count is less than the mean temperature fault count, or if the GPU power consumption fault standard deviation is greater than the GPU power consumption standard deviation threshold, the comparing determines that the GPU is not to malfunction.

[0012] Further alternatively, the GPU status parameters collected from the GPU node include a GPU model, power consumption, and usage status. The mean status fault parameters are obtained by use of a pre-configured statistical model through the following steps. First, it is determined,

based on the usage status, whether the GPU malfunctions. Second, in response to a determination that the GPU does not malfunction, based upon the GPU model, the power consumption collected from the GPU node is stored in an information storage space corresponding to the GPU model. In response to a determination that the GPU does malfunction, based on the GPU model, stored power consumption are obtained from an information storage space corresponding to the GPU model. Based on the power consumption collected from the GPU node and the stored power consumption obtained from the storage space corresponding to the GPU model, statistics are performed to compute an arithmetic mean power consumption fault count and a power consumption fault standard deviation, by use of a pre-configured power consumption statistical model.

[0013] Still further alternatively, the pre-configured power consumption statistical model includes a mean power consumption fault count model configured to compute an arithmetic mean based on the GPU power consumption collected from the GPU node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model. The pre-configured power consumption statistical model also includes a power consumption fault standard deviation model configured to compute a power consumption standard deviation based on the GPU power consumption collected from the GPU node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model, and the mean power consumption fault count.

[0014] Alternatively, when the status parameter is usage duration, the comparing the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction includes comparing the obtained GPU usage duration with a mean fault usage duration. The comparing further determines that, if the GPU usage duration is greater than the mean fault usage duration, the GPU is to malfunction. Otherwise, if the GPU usage duration is less than the mean fault usage duration, the comparing determines that the GPU is not to malfunction.

[0015] Further alternatively, the GPU status parameters collected from the GPU node include a GPU model, usage duration, and usage status. The mean status fault parameters are obtained by use of a pre-configured statistical model through the following steps. First, it is determined, based on the usage status, whether the GPU malfunctions. Second, in response to a determination that the GPU does not malfunction, based upon the GPU model, the usage duration collected from the GPU node is stored in an information storage space corresponding to the GPU model. In response to a determination that the GPU does malfunction, based on the GPU model, the stored usage duration is obtained from an information storage space corresponding to the GPU model, and based on the

usage duration collected from the GPU node and the stored usage duration obtained from the storage space corresponding to the GPU model, statistics are performed to compute an arithmetic mean fault usage duration, by use of a pre-configured usage duration statistical model.

[0016] Still further alternatively, the pre-configured usage duration statistical model includes a mean fault usage duration model configured to compute an arithmetic mean based on the GPU usage duration collected from the GPU node and the stored GPU usage duration obtained from the information storage space corresponding to the GPU model.

[0017] According to another exemplary embodiment of the present disclosure, an apparatus for predicting GPU malfunctions includes an installation module, a collecting module and a processing module. The installation module is configured to install a daemon program at a GPU node, the daemon program periodically collecting GPU status parameters corresponding to the GPU node at a pre-determined time period. The collecting module is configured to obtain the GPU status parameters from the GPU node. The processing module is configured to compare the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, where the mean status fault parameters are obtained by use of a pre-configured statistical model.

[0018] Alternatively, when the status parameter is temperature, the processing module includes a first statistical module, a first comparison module, a first determination module and a second determination module. The first statistical module is configured to perform statistics to obtain a temperature count of a number of times the GPU incurs a temperature greater than a pre-determined temperature threshold and a GPU temperature standard deviation threshold for the GPU. The first comparison module is configured to compare the obtained temperature count with a mean temperature fault count and to compare a GPU temperature fault standard deviation with the GPU temperature standard deviation threshold. The first determination module is configured to, if the temperature count is greater than the mean temperature fault count, and if the GPU temperature fault standard deviation is less than the GPU temperature standard deviation threshold, determine that the GPU is to malfunction. The second determination module is configured to, if the temperature count is less than the mean temperature fault count, or if the GPU temperature fault standard deviation is greater than the GPU temperature standard deviation threshold, determine that the GPU is not to malfunction.

[0019] Further alternatively, after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period. The collecting module includes a

first collecting module configured to collect from the GPU node a GPU model, temperature and usage status. The processing module further includes a first decision module, a first storing module and a first computing module. The first decision module is configured to decide, based on the usage status, whether the GPU malfunctions. The first storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the temperature collected from the GPU node in an information storage space corresponding to the GPU model. The first computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtain temperatures stored in an information storage space corresponding to the GPU model, and based on the temperature collected from the GPU node and the stored temperatures obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean temperature fault count and a temperature fault standard deviation, by use of a pre-configured temperature statistical model.

[0020] Still further alternatively, the pre-configured temperature statistical model includes a mean temperature fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model. The pre-configured temperature statistical model also includes a temperature fault standard deviation model configured to compute a standard deviation based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, and the mean temperature fault count.

[0021] Yet still alternatively, when the status parameter is power consumption, the processing module includes a second statistical module, a second comparison module, a third determination module and a fourth determination module. The second statistical module is configured to perform statistics to obtain a power consumption count of a number of times the GPU incurs power consumption greater than a pre-determined power consumption threshold and a GPU power consumption standard deviation threshold for the GPU. The second comparison module is configured to compare the obtained power consumption count with a mean power consumption fault count and to compare a GPU power consumption fault standard deviation with the GPU power consumption standard deviation threshold. The third determination module is configured to, if the power consumption count is greater than the mean power consumption fault count, and if the GPU power consumption fault standard deviation is less than the GPU power consumption standard deviation threshold, determine that the GPU is to malfunction. The fourth determination module is

configured to, if the power consumption count is less than the mean power consumption fault count, or if the GPU power consumption fault standard deviation is greater than the GPU power consumption standard deviation threshold, determine that the GPU is not to malfunction.

[0022] Yet still further alternatively, after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period. The collecting module includes a second collecting module configured to collect from the GPU node a GPU model, power consumption and usage status. The processing module further includes a second decision module, a second storing module and a second computing module. The second decision module is configured to decide, based on the usage status, whether the GPU malfunctions. The second storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the power consumption collected from the GPU node in an information storage space corresponding to the GPU model. The second computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtain power consumption stored in an information storage space corresponding to the GPU model, and based on the power consumption collected from the GPU node and the stored power consumption obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean power consumption fault count and a power consumption fault standard deviation, by use of a pre-configured power consumption statistical model.

[0023] Still further alternatively, the pre-configured power consumption statistical model includes a mean power consumption fault count model configured to compute an arithmetic mean based on the GPU power consumption collected from the GPU node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model. The pre-configured power consumption statistical model also includes a power consumption fault standard deviation model configured to compute a standard deviation based on the GPU power consumption collected from the GPU node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model, and the mean power consumption fault count.

[0024] Alternatively, when the status parameter is usage duration, the processing module includes a third comparison module, a fifth determination module and a sixth determination module. The third comparison module is configured to compare the obtained GPU usage duration with a mean fault usage duration. The fifth determination module is configured to, if the GPU usage

duration is greater than the mean fault usage duration, determine the GPU is to malfunction. The sixth determination module is configured to, if the GPU usage duration is less than the mean fault usage duration, determine the GPU is not to malfunction.

[0025] Yet still further alternatively, after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period. The collecting module includes a third collecting module configured to collect from the GPU node a GPU model, usage duration and usage status. The processing module further includes a third decision module, a third storing module and a third computing module. The third decision module is configured to decide, based on the usage status, whether the GPU malfunctions. The third storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the usage duration collected from the GPU node in an information storage space corresponding to the GPU model. The third computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtain usage duration stored in an information storage space corresponding to the GPU model, and based on the usage duration collected from the GPU node and the stored usage duration obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean fault usage duration, by use of a pre-configured usage duration statistical model.

[0026] Still further alternatively, the pre-configured usage duration statistical model includes a mean fault usage duration model configured to compute an arithmetic mean based on the GPU usage duration collected from the GPU node and the stored GPU usage duration obtained from the information storage space corresponding to the GPU model.

[0027] Compared to the current technology, the present disclosure provides for the following technical effects. 1) Prior to a GPU enters a malfunction state, status parameters and mean status fault parameters can be utilized to determine whether a GPU is about to malfunction such that a malfunction state of the GPU can be predicted. Consequently, prior to the GPU enters a malfunction state, the GPU can be replaced, or the programs executing on the GPU can be migrated to other GPUs for execution, without affecting the normal business operations.

[0028] 2) Prior to the GPU enters a malfunction state, the GPU temperature count and mean temperature fault count, power consumption count and mean power consumption fault count, or usage duration and mean fault usage duration, etc. can be utilized to determine whether the GPU is about to malfunction such that a malfunction state of the GPU can be predicted. Consequently, prior

to the GPU enters a malfunction state, the GPU can be replaced, or the programs executing on the GPU can be migrated to other GPUs for execution, without affecting the normal business operations.

[0029] 3) When the GPU temperature count and mean temperature fault count, temperature standard deviation threshold and temperature fault standard deviation are utilized to determine whether the GPU is about to malfunction, based on the GPU model, the mean temperature fault count and the temperature fault standard deviation corresponding to the GPU model can be obtained to enhance the GPU malfunction prediction accuracy.

[0030] 4) When the GPU power consumption count and mean power consumption fault count, power consumption standard deviation threshold and power consumption fault standard deviation are utilized to determine whether the GPU is about to malfunction, based on the GPU model, the mean power consumption fault count and the power consumption fault standard deviation corresponding to the GPU model can be obtained to enhance the GPU malfunction prediction accuracy.

[0031] 5) When the GPU usage duration and mean fault usage duration are utilized to determine whether the GPU is about to malfunction, based on the GPU model, the mean fault usage duration corresponding to the GPU model can be obtained to enhance the GPU malfunction prediction accuracy.

[0032] It should be appreciated by one having ordinary skills of the art that embodiments of the present disclosure do not need to implement or achieve all the above described technical effects.

DESCRIPTION OF THE DRAWINGS

[0033] The accompanying drawings, which are incorporated in and form a part of this specification and in which like numerals depict like elements, illustrate embodiments of the present disclosure and, together with the description, serve to explain the principles of the disclosure.

[0034] FIG. 1 is a flow chart of a first exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0035] FIG. 2 is a flow chart of a second exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0036] FIG. 3 is a flow chart of a third exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0037] FIG. 4 is a flow chart of a fourth exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0038] FIG. 5 is a flow chart of a fifth exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0039] FIG. 6 is a flow chart of a sixth exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure;

[0040] FIG. 7 is a flow chart of a seventh exemplary method of predicting GPU malfunctions in accordance with an embodiment of the present disclosure; and

[0041] FIG. 8 is a block diagram of an exemplary apparatus for predicting GPU malfunctions in accordance with an embodiment of the present disclosure.

DETAILED DESCRIPTION

[0042] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present disclosure. However, it will become obvious to those skilled in the art that the present disclosure may be practiced without these specific details. The descriptions and representations herein are the common means used by those experienced or skilled in the art to most effectively convey the substance of their work to others skilled in the art. In other instances, well-known methods, procedures, components, and circuitry have not been described in detail to avoid unnecessarily obscuring aspects of the present disclosure.

[0043] Reference herein to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment can be included in at least one embodiment of the disclosure. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Used herein, the terms “upper”, “lower”, “top”, “bottom”, “middle”, “upwards”, and “downwards” are intended to provide relative positions for the purposes of description, and are not intended to designate an absolute frame of reference. Further, the order of blocks in process flowcharts or diagrams representing one or more embodiments of the disclosure do not inherently indicate any particular order nor imply any limitations in the disclosure.

[0044] Embodiments of the present disclosure are discussed herein with reference to FIGS. 1-8. However, those skilled in the art will readily appreciate that the detailed description given herein

with respect to these figures is for explanatory purposes as the disclosure extends beyond these limited embodiments.

[0045] Failure, failure conditions, faults, fault conditions, errors, malfunctions, malfunction states and malfunction conditions refer herein to any problem, error, fault, out of order or tolerable condition. Status fault parameters refer herein to the values of the status parameter associated with the faulty state caused by the status parameter. Temperature faults refer herein to malfunction states caused by temperature conditions. Power consumption faults refer herein to malfunction states caused by conditions related to power consumption. Usage duration faults refer herein to malfunction states caused by usages for certain durations.

[0046] Referring to FIG. 1, a flow chart of a first exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. The method 100 starts in step S101, where a daemon program is installed at a GPU node. The daemon program periodically collects status parameters corresponding to the GPU node accordingly to a pre-determined time period. In particular, a daemon program (e.g., aswp agent) installed for a GPU node periodically collects status parameters such as a GPU model, temperature, power consumption, usage duration and usage status, etc. corresponding to the GPU node, according to a pre-determined time period.

[0047] In step S102, the GPU status parameters are obtained from the GPU node. In step S103, the obtained GPU status parameters are compared to the mean status fault parameters, the result of which is utilized to determine whether the GPU is about to malfunction. The mean status fault parameters refer herein to the mean value of the status parameters that are associated with the fault conditions. The mean status fault parameters are obtained through pre-configured statistical models.

[0048] In particular, the mean status fault parameters can be pre-configured according to the prior practical experience. Alternatively, the mean status fault parameters can also be computed by use of the pre-configured statistical models to perform statistics over the status parameter data of the already faulty GPUs. The present embodiment focuses on how to utilize the pre-configured statistical models to generate mean status fault parameters based on the status parameter data of the GPUs that have already incurred faults or errors.

[0049] Referring to FIG. 2, a flow chart of a second exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In embodiments, the status parameter is temperature. Step S103, where the GPU status parameters are compared to the mean status fault parameters, the result of which is utilized to determine whether

the GPU is about to malfunction, can include steps S201 through S204. Method 200 starts in step S201, where the number of times the GPU incurs a temperature greater than a pre-configured temperature threshold (“temperature count”) is obtained. Further, the GPU temperature standard deviation threshold is also obtained.

[0050] In particular, it can be determined whether the GPU is to enter a malfunction state by use of only the GPU temperature count and the mean temperature fault count. In order to increase the accuracy of prediction, a temperature fault standard deviation and a temperature standard deviation threshold can be utilized in addition to predict GPU malfunctions. In other words, based upon the GPU temperature count, mean temperature fault count, as well as temperature fault standard deviation and temperature standard deviation threshold, it can be determined whether the GPU is about to malfunction.

[0051] Here, the pre-configured temperature threshold can be configured based on the information such as the maximal temperature the GPU can incur, or the like. The GPU temperature standard deviation threshold can be configured according to the practical use experience, or be configured based upon multiple experiments or the like. The pre-configured temperature threshold, GPU temperature standard deviation threshold can be stored in a one-to-one mapping relationship corresponding to the GPU in a correspondent storage space for retrieval at a time when such information is needed.

[0052] At step S202, the GPU temperature count is compared to the mean temperature fault count, the temperature fault standard deviation is compared to the temperature standard deviation threshold. If the GPU temperature count is greater than the mean temperature fault count, and the temperature fault standard deviation is less than the temperature standard deviation threshold, method 200 proceeds to step S203. If the GPU temperature count is less than the mean temperature fault count, or the temperature fault standard deviation is greater than the temperature standard deviation threshold, method 200 proceeds to step S204.

[0053] In particular, the mean temperature fault count can be configured based on the practical use experience. It can also be computed by use of statistical models to perform statistics of the status parameter data of the already faulty GPUs. For example, statistics can be performed for each faulty GPU to compute the temperature counts, i.e., the number of times the GPU incurs temperatures exceeding the pre-configured temperature threshold. The mean temperature fault count can be obtained by computing an arithmetic mean of all the GPU temperature counts.

[0054] In step S203, method 200 determines that the GPU is to malfunction and concludes. In particular, determining that the GPU is to malfunction is to predict that the GPU will enter into a malfunction state in the future. Therefore, before the GPU indeed enters into a malfunction state, the GPU can be replaced in advance, or the programs executing on the GPU can be migrated to the other GPUs for execution, without affecting the normal business operations. Further, after it is determined that the GPU is to malfunction, alerts can be provided by use of sound alerts (e.g., beeps, etc.), phrase alerts (e.g., hints, etc.), voice alerts, or the like.

[0055] In step S204, method 200 determines that the GPU is not malfunction and proceeds to step S102 of FIG. 1. In particular, when it is determined that the GPU is not to malfunction, step S102 is executed such that to continue to monitor the status parameters of the GPU.

[0056] Referring to FIG. 3, a flow chart of a third exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In particular, in this embodiment, after step S101 where the daemon programs are installed in the GPU nodes, the daemon programs collect periodically from the GPU node the GPU model and GPU status parameters corresponding to the GPU node, according to a pre-determined time period. Step S102, where the GPU status parameters are obtained from the GPU node, includes step S301, where the GPU model, temperature and usage status are obtained.

[0057] Step S103, where the mean status fault parameters are obtained by use of statistical models includes steps S302 through S304. In step S302, it is determined whether the GPU malfunctions based on the GPU usage status. If it is determined that the GPU does not malfunction, along the NO path, method 300 executes step S303; otherwise if it is determined that the GPU does malfunction, along the YES path, method 300 executes step S304.

[0058] In step S303, based on the GPU model, the GPU temperature obtained from the GPU node is stored in an information storage space corresponding to the GPU model, after which method 300 executes step S301.

[0059] In step S304, based on the GPU model, the GPU temperatures are obtained from the information storage space corresponding to the GPU model. Based on the GPU temperature obtained from the GPU node, as well as the stored GPU temperatures which are obtained from the information storage space corresponding to the GPU model, by use of a pre-configured temperature statistical model, statistics are performed separately to obtain a mean temperature fault count and temperature fault standard deviation. Afterwards, method 300 executes step S301.

[0060] The pre-configured temperature statistical model includes a mean temperature fault count model and a temperature fault standard deviation model. The mean temperature fault count model is configured to obtain the GPU temperatures from the GPU node and the stored GPU temperatures from the information storage space corresponding to the GPU model such that to compute an arithmetic mean of the temperature counts.

[0061] The temperature fault standard deviation model is configured to obtain the GPU temperatures from the GPU node and the stored GPU temperatures from the information storage space corresponding to the GPU model such that, together with the mean temperature count, to compute a temperature fault standard deviation.

[0062] In particular, in a hypothetical condition where a GPU of a particular GPU model functions normally, the GPU reaches a pre-configured temperature threshold T. When the GPU of such GPU model malfunctions, there are already n-1 number of the GPUs of the same GPU model have incurred faults, each faulty GPU's respective temperature count of a temperature greater than the temperature T is NT1, NT2, . . . , NTn-1. Based upon the temperature obtained from the GPU node, and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, it can be determined that the GPU has incurred a temperature greater than T for a number of NTn times. For this particular GPU's model, a mean temperature fault count NT can be computed with the formula: $NT = (NT1 + NT2 + \dots + NTn) / n$; and a temperature fault standard

deviation $\sigma(NT)$ can be computed with the formula: $\sigma(NT) = \sqrt{\frac{1}{n} (\sum_{i=1}^n (NT_i - NT)^2)}$.

[0063] In particular, with the computed mean GPU temperature fault count and the temperature fault standard deviation, method 300 continues to execute step S301 such that when there is a new occurrence of GPU fault, a new mean GPU temperature fault count and temperature fault standard deviation can be computed. Therefore, the mean GPU temperature fault count and temperature fault standard deviation can be updated constantly so that the mean GPU temperature fault count and temperature fault standard deviation can be computed with more accuracy, and the prediction of GPU malfunctions can be performed with more accuracy accordingly.

[0064] In particular, for GPUs of different GPU models, the status parameters such as temperature corresponding to a malfunction state can vary in a large degree. Therefore, based on the GPU models, a GPU mean temperature fault count and temperature fault standard deviation are computed for each GPU model. Thus, the determining of whether a GPU is to malfunction based upon the GPU temperature count and mean temperature fault count, as well as the temperature fault

standard deviation and temperature standard deviation threshold, can be performed with enhanced accuracy.

[0065] Referring to FIG. 4, a flow chart of a fourth exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In some embodiments, the status parameter is power consumption. Step S103 of comparing the GPU status parameters with the mean status parameters such that to generate a result for determining whether the GPU is about to malfunction includes steps S401 through S404.

[0066] In step S401, statistics are performed over the power consumption count, i.e., the number of times the GPU incurs power consumption that is greater than a pre-determined power consumption threshold. The GPU power consumption standard deviation threshold is also obtained.

[0067] In particular, based upon the GPU power consumption count and the mean power consumption fault count only, it can be determined whether the GPU is about to enter in a malfunction state. However, in order to increase the accuracy of predicting GPU malfunctions, the power consumption fault standard deviation and power consumption standard deviation threshold can also be utilized in addition to predict GPU malfunctions. In other words, based upon the GPU power consumption count, mean power consumption fault count, as well as power consumption fault standard deviation and power consumption standard deviation threshold, it can be determined whether the GPU is about to malfunction.

[0068] Here, the pre-determined power consumption threshold can be configured according to the information such as the maximal power consumption the GPU can possibly incur or the like. The GPU power consumption standard deviation threshold can be configured based upon the practical use experience, or upon multiple experiments or the like. The pre-configured power consumption threshold, GPU power consumption standard deviation threshold can be stored in a one-to-one mapping relationship corresponding to the GPU in a correspondent storage space for retrieval at a time when such information is needed.

[0069] In step S402, the GPU power consumption count is compared to the mean power consumption fault count, the power consumption fault standard deviation is compared to the power consumption standard deviation threshold. If the GPU power consumption count is greater than the mean power consumption fault count, and the power consumption fault standard deviation is less than the power consumption standard deviation threshold, method 400 executes steps S403. If the GPU power consumption count is less than the mean power consumption fault count, or the power

consumption fault standard deviation is greater than the power consumption standard deviation threshold, method 400 executes steps S404.

[0070] In particular, based on the practical use experience the mean power consumption fault count can be configured. It can also be configured by performing statistics over the GPUs that have already incurred faults to obtain the mean power consumption fault count. For example, statistics can be performed for each faulty GPU to compute the power consumption counts, i.e., the number of times the GPU incurs power consumption exceeding the pre-configured power consumption threshold. The mean power consumption fault count can be obtained by computing an arithmetic mean of all the GPU power consumption counts.

[0071] In step S403, method 400 determines that the GPU is to malfunction and concludes. In particular, determining that the GPU is to malfunction is to predict that the GPU is likely to enter a malfunction state. Therefore, before the GPU malfunctions, the GPU can be replaced in advance, or the programs executing on the GPU can be migrated to the other GPUs for execution, without affecting the normal business operations. Further, after it is determined that the GPU is to malfunction, alerts can be provided by use of sound alerts (e.g., beeps, etc.), phrase alerts (e.g., hints, etc.), voice alerts, or the like.

[0072] In step S404, it is determined that the GPU is not to malfunction, method 400 proceeds to step S102. In particular, if the GPU is determined not to malfunction, step S102 is executed such that to continue to monitor the status parameters of the GPU.

[0073] Referring to FIG. 5, a flow chart of a fifth exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In particular, in this embodiment, after step S101 where the daemon programs are installed in the GPU nodes, the daemon programs collect periodically from the GPU node the GPU model and GPU status parameters corresponding to the GPU node, according to a pre-determined time period. Step S102, where the GPU status parameters are obtained from the GPU node, includes step S501, where the GPU model, power consumption and usage status are obtained.

[0074] Step S103, where the mean status fault parameters are obtained by use of statistical models includes steps S502 through S504. In step S502, it is determined whether the GPU malfunctions based on the GPU usage status. If it is determined that GPU does not malfunction, along the NO path, method 500 executes step S503; otherwise if it is determined that the GPU does malfunction, along the YES path, method 500 executes step S504.

[0075] In step S503, based on the GPU model, the GPU power consumption obtained from the GPU node is stored in an information storage space corresponding to the GPU model, after which method 500 executes step S501.

[0076] In step S504, based on the GPU model, the GPU power consumption are obtained from the information storage space corresponding to the GPU model. Based on the GPU power consumption obtained from the GPU node, as well as the stored GPU power consumption which is obtained from the information storage space corresponding to the GPU model, by use of a pre-configured power consumption statistical model, statistics are performed separately to obtain a mean power consumption fault count and power consumption fault standard deviation. Afterwards, method 500 executes step S501.

[0077] The pre-configured power consumption statistical model includes a mean power consumption fault count model and a power consumption fault standard deviation model. The mean power consumption fault count model is configured to obtain the GPU power consumption from the GPU node and the stored GPU power consumption from the information storage space corresponding to the GPU model such that to compute an arithmetic mean of the power consumption counts.

[0078] The power consumption fault standard deviation model is configured to obtain the GPU power consumption from the GPU node and the stored GPU power consumption from the information storage space corresponding to the GPU model such that, together with the mean power consumption count, to compute a power consumption fault standard deviation.

[0079] In particular, in a hypothetical condition where a GPU of a particular GPU model functions normally, the GPU reaches a pre-configured power consumption threshold W. When the GPU of such GPU model malfunctions, there are already n-1 number of the GPUs of the same GPU model have incurred faults, each faulty GPU's respective power consumption count of power consumption greater than the power consumption W is NW1, NW2, . . . , NWn-1. Based upon the power consumption obtained from the GPU node, and the stored GPU power consumption s obtained from the information storage space corresponding to the GPU model, it can be determined that the GPU has incurred power consumption greater than W for a number of NWn times. For this particular GPU's model, a mean power consumption fault count NW can be computed with the formula: $NW = (NW1 + NW2 + \dots + NWn)/n$; and a power consumption fault standard deviation

$\sigma(NW)$ can be computed with the formula: $\sigma(NW) = \sqrt{\frac{1}{n}(\sum_{i=1}^n (NW_i - NW))^2}$.

[0080] In particular, with the computed mean GPU power consumption fault count and the power consumption fault standard deviation, method 500 continues to execute step S501 such that when there is a new occurrence of GPU fault, a new mean GPU power consumption fault count and power consumption fault standard deviation can be computed. Therefore, the mean GPU power consumption fault count and power consumption fault standard deviation can be updated constantly so that the mean GPU power consumption fault count and power consumption fault standard deviation can be computed with more accuracy, and the prediction of GPU malfunctions can be performed with more accuracy accordingly.

[0081] In particular, for GPUs of different GPU models, the status parameters such as power consumption corresponding to a malfunction state can vary in a large degree. Therefore, based on the GPU models, a GPU mean power consumption fault count and power consumption fault standard deviation are computed for each GPU model. Thus, the determining of whether a GPU is to malfunction based upon the GPU power consumption count and mean power consumption fault count, as well as the power consumption fault standard deviation and power consumption standard deviation threshold, can be performed with enhanced accuracy.

[0082] Referring to FIG. 6, a flow chart of a sixth exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In some embodiments, the status parameter is usage duration. Step S103, where the GPU status parameters are compared to the mean status fault parameters, the result of which is utilized to determine whether the GPU is about to malfunction, can include steps S601 through S603. Method 600 starts in step S601, where the GPU usage duration is compared to a mean fault usage duration. If the GPU usage duration is greater than the mean fault usage duration, method 600 proceeds to step S602. If the GPU usage duration is less than the mean fault usage duration, method 600 proceeds to step S603.

[0083] In step S602, method 600 determines that the GPU is to enter a malfunction state and concludes. In particular, determining that the GPU is to malfunction is to predict that the GPU is likely to enter a malfunction state in the future. Therefore, before the GPU malfunctions, the GPU can be replaced in advance, or the programs executing on the GPU can be migrated to the other GPUs for execution, without affecting the normal business operations. Further, after it is determined that the GPU is to malfunction, alerts can be provided by use of sound alerts (e.g., beeps, etc.), phrase alerts (e.g., hints, etc.), voice alerts, or the like.

[0084] In step S603, method 600 determines that the GPU is not to malfunction and proceeds to step S102. In particular, when it is determined that the GPU is not to malfunction, step S102 is executed such that to continue to monitor the status parameters of the GPU.

[0085] Referring to FIG. 7, a flow chart of a seventh exemplary method of predicting GPU malfunctions is shown in accordance with an embodiment of the present disclosure. In particular, in this embodiment, after step S101 where the daemon programs are installed in the GPU nodes, the daemon programs collect periodically from the GPU node the GPU model and GPU status parameters corresponding to the GPU node, according to a pre-determined time period. Step S102, where the GPU status parameters are obtained from the GPU node, includes step S701, where the GPU model, usage duration and usage status are obtained.

[0086] The obtaining of the mean status fault parameters by use of statistical models includes steps S702 through S704. In step S702, it is determined whether the GPU malfunctions based on the GPU usage status. If it is determined that GPU does not malfunction, along the NO path, method 700 executes step S703; otherwise if it is determined that the GPU does malfunction, along the YES path, method 700 executes step S704.

[0087] In step S703, based on the GPU model, the GPU usage duration obtained from the GPU node is stored in an information storage space corresponding to the GPU model, after which method 700 executes step S701.

[0088] In step S704, based on the GPU model, the GPU usage duration are obtained from the information storage space corresponding to the GPU model. Based on the GPU usage duration obtained from the GPU node, as well as the stored GPU usage duration which is obtained from the information storage space corresponding to the GPU model, by use of a pre-configured usage duration statistical model, statistics are performed obtain a mean fault usage duration. Afterwards, method 700 executes step S701.

[0089] Here, the pre-configured usage duration statistical model includes a mean fault usage duration model. The mean fault usage duration model is configured to obtain the GPU usage duration from the GPU node and the stored GPU usage duration from the information storage space corresponding to the GPU model such that to compute an arithmetic mean of the usage duration.

[0090] In particular, in a hypothetical condition where a GPU of a particular GPU model malfunctions, there are already n-1 number of the GPUs of the same GPU model have incurred faults, each faulty GPU's respective usage duration is NS1, NS2, . . . , NSn-1. Based upon the GPU usage duration obtained from the GPU node, and the stored GPU usage duration obtained from the

information storage space corresponding to the GPU model, it can be determined that, for this particular GPU's model, a mean fault usage duration NS can be computed with the formula: $NS = (NS1+NS2+...+NSn)/n$.

[0091] In particular, with the computed mean GPU fault usage duration, method 700 continues to execute step S701 such that when there is a new occurrence of GPU fault, a new mean GPU fault usage duration can be computed. Therefore, the mean GPU fault usage duration can be updated constantly so that the mean GPU fault usage duration can be computed with more accuracy, and the prediction of GPU malfunctions can be performed with more accuracy accordingly.

[0092] In particular, for GPUs of different GPU models, the status parameters such as usage duration corresponding to a malfunction state can vary in a large degree. Therefore, based on the GPU models, a GPU mean fault usage duration is computed for each GPU model. Thus, the determining of whether a GPU is to malfunction, based upon the GPU usage duration and the mean GPU fault usage duration, can be performed with enhanced accuracy.

[0093] In accordance with embodiments of the present disclosure, prior to a GPU malfunctions, it can be determined based on the status parameters and mean status fault parameters whether the GPU is to enter a malfunction state. Therefore, before the GPU malfunctions, the GPU can be replaced, or the programs executing on the GPU can be migrated to the other GPUs for execution, without affecting the normal business operations.

[0094] Before a GPU incurs faults, based on the GPU temperature count and a mean temperature fault count, power consumption count and a mean power consumption fault count, or usage duration and a mean fault usage duration, etc. to determined whether the GPU is to generate faults. With such prediction in advance whether the GPU is to malfunction, before the GPU faults, the GPU can be replaced in advance, or the programs executing on the GPU can be migrated to the other GPUs for execution, without affecting the normal business operations.

[0095] When the GPU temperature count and mean temperature fault count, temperature fault standard deviation and temperature standard deviation threshold are utilized to determine whether the GPU is to malfunction, based on the GPU model, the temperature mean fault count and the temperature fault standard deviation corresponding to the GPU model can be obtained such that to enhance the prediction accuracy.

[0096] Similarly, when the GPU power consumption count and mean power consumption fault count, power consumption fault standard deviation and power consumption standard deviation

threshold are utilized to determine whether the GPU is to malfunction, based on the GPU model, the power consumption mean fault count and the power consumption fault standard deviation corresponding to the GPU model can be obtained such that to enhance the prediction accuracy.

[0097] Further similarly, when the GPU usage duration is utilized to determine whether the GPU is to malfunction, based on the GPU model, the mean GPU usage fault duration corresponding to the GPU model can also be obtained such that to enhance the prediction accuracy.

[0098] Referring to FIG. 8, a block diagram of an exemplary apparatus for predicting GPU malfunction is shown in accordance with an embodiment of the present disclosure. Apparatus 800 includes an installing module 801, a collecting module 802 and a processing module 803. The installing module 801 is configured to install a daemon program at a GPU node, the daemon program periodically collecting GPU status parameters corresponding to the GPU node at a pre-determined time period.

[0099] The collecting module 802 is configured to obtain the GPU status parameters from the GPU node. The processing module 803 is configured to compare the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, where the mean status fault parameters are obtained by use of pre-configured statistical models.

[00100] Further, when the status parameter is temperature, the processing module 803 includes a first statistical module, a first comparison module, a first determination module and a second determination module. The first statistical module is configured to perform statistics to obtain a temperature count of a number of times the GPU incurs a temperature greater than a pre-determined temperature threshold and a GPU temperature standard deviation threshold for the GPU. The first comparison module is configured to compare the obtained temperature count with a mean temperature fault count and to compare a GPU temperature standard deviation with the GPU temperature standard deviation threshold.

[00101] The first determination module is configured to, if the temperature count is greater than the mean temperature fault count, and if the GPU temperature standard deviation is less than the GPU temperature standard deviation threshold, determine that the GPU is to malfunction. The second determination module is configured to, if the temperature count is less than the mean temperature fault count, or if the GPU temperature standard deviation is greater than the GPU temperature standard deviation threshold, determine that the GPU is not to malfunction.

[00102] After the installation module 801 installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to

the GPU node at a pre-determined time period. Correspondingly, the collecting module 802 includes a first collecting module configured to collect the GPU model, temperature and usage status from the GPU node. Also correspondingly, the processing module 803 further includes a first decision module, a first storing module and a first computing module. The first decision module is configured to decide, based on the usage status, whether the GPU malfunctions. The first storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the temperature collected from the GPU node in an information storage space corresponding to the GPU model.

[00103] The first computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtaining temperatures stored in an information storage space corresponding to the GPU model, and based on the temperature collected from the GPU node and the stored temperatures obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean temperature fault count and a GPU temperatures fault standard deviation, by use of a pre-configured temperature statistical model.

[00104] Further, the pre-configured temperature statistical model includes a mean temperature fault count model and a temperature fault standard deviation model. The mean temperature fault count model is configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model.

[00105] The temperature fault standard deviation model is configured to compute a temperature fault standard deviation based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, together with the mean temperature fault count.

[00106] Furthermore, when the status parameter is power consumption, the processing module 803 includes a second statistical module, a second comparison module, a third determination module and a fourth determination module. The second statistical module is configured to perform statistics to obtain a power consumption count of the number of times the GPU incurs power consumption greater than a pre-determined power consumption threshold and a GPU power consumption standard deviation threshold for the GPU. The second comparison module is configured to compare the obtained power consumption count with a mean power consumption fault count and to compare a GPU power consumption fault standard deviation with the GPU power consumption standard deviation threshold.

[00107] The third determination module is configured to, if the power consumption count is greater than the mean power consumption fault count, and if the GPU power consumption fault standard deviation is less than the GPU power consumption standard deviation threshold, determine that the GPU is to malfunction. The fourth determination module is configured to, if the power consumption count is less than the mean power consumption fault count, and or the GPU power consumption fault standard deviation is greater than the GPU power consumption standard deviation threshold, determine that the GPU is not to malfunction.

[00108] After the installation module 801 installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period. Correspondingly, the collecting module 802 includes a second collecting module configured to collect from the GPU node a GPU model, power consumption, and usage status. Also correspondingly, the processing module 803 further includes a second decision module, a second storing module and a second computing module.

[00109] The second decision module is configured to, based on the usage status, decide whether the GPU malfunctions. The second storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the power consumption collected from the GPU node in an information storage space corresponding to the GPU model.

[00110] The second computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtain power consumption stored in an information storage space corresponding to the GPU model. Based on the power consumption collected from the GPU node and the stored power consumption obtained from the storage space corresponding to the GPU model, the second computing module is configured to perform statistics to compute an arithmetic mean power consumption fault count and a power consumption fault standard deviation, by use of a pre-configured power consumption statistical model.

[00111] The pre-configured power consumption statistical model includes a mean power consumption fault count model and a power consumption fault standard deviation model. The mean power consumption fault count model is configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model.

[00112] The power consumption fault standard deviation model is configured to compute a power consumption fault standard deviation based on the GPU power consumption collected from the GPU

node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model, together with the mean power consumption fault count.

[00113] When the status parameter is usage duration, the processing module 803 includes a third comparison module, a fifth determination module and a sixth determination module. The third comparison module is configured to compare the obtained GPU usage duration with a mean fault usage duration. The fifth determination module is configured to, if the GPU usage duration is greater than the mean fault usage duration, determine the GPU is to malfunction. The sixth determination module is configured to, if the GPU usage duration is less than the mean fault usage duration, determine the GPU is not to malfunction.

[00114] After the installation module 801 installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period. Correspondingly, the collecting module 802 includes a third collecting module configured to collect from the GPU node a GPU model, usage duration, and usage status. Also correspondingly, the processing module 803 further includes a third decision module, a third storing module and a third computing module. The third decision module is configured to, based on the usage status, decide whether the GPU malfunctions. The third storing module is configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the usage duration collected from the GPU node in an information storage space corresponding to the GPU model.

[00115] The third computing module is configured to, in response to a determination that the GPU does malfunction, based on the GPU model, obtain usage duration stored in an information storage space corresponding to the GPU model. Based on the usage duration collected from the GPU node and the stored usage duration obtained from the storage space corresponding to the GPU model, the third computing module is configured to perform statistics to compute an arithmetic mean fault usage duration, by use of a pre-configured usage duration statistical model.

[00116] The pre-configured usage duration statistical model includes a mean fault usage duration model configured to compute an arithmetic mean based on the GPU usage duration collected from the GPU node and the stored GPU usage duration obtained from the information storage space corresponding to the GPU model.

[00117] Embodiments of the present disclosure can be implemented using software, hardware, firmware, and/or the combinations thereof. Regardless of being implemented using software, hardware, firmware or the combinations thereof, instruction code can be stored in any kind of

computer readable media (for example, permanent or modifiable, volatile or non-volatile, solid or non-solid, fixed or changeable medium, etc.). Similarly, such medium can be implemented using, for example, programmable array logic (PAL), random access memory (RAM), programmable read only memory (PROM), read only memory (ROM), electrically erasable programmable ROM (EEPROM), magnetic storage, optical storage, digital versatile disc (DVD), or the like.

[00118] It is necessary to point out that, modules or blocks described by embodiments of the present disclosures are logical modules or logical blocks. Physically, a logical module or logical block can be a physical module or a physical block, a part of a physical module or a physical block, or the combinations of more than one physical modules or physical blocks. Physical implementation of those logical module or logical blocks is not of essence. The realized functionalities realized by the modules, blocks and the combinations thereof are key to solving the problems addressed by the present disclosure. Further, in order to disclose the novelties of the present disclosure, the above described embodiments do not disclose about those modules or blocks not too related to solving the problems addressed by the present disclosure, which does not mean that the above described embodiments cannot include other modules or blocks.

[00119] It is also necessary to point out that, in the claims and specification of the present disclosure, terms such as first and second only are for distinguishing an embodiment or an operation from another embodiment or operation. It does not require or imply that those embodiments or operations having any such real relationship or order. Further, as used herein, the terms “comprising,” “including,” or any other variation intended to cover a non-exclusive inclusion, such that a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. Absent further limitation, elements recited by the phrase “comprising a” does not exclude a process, method, article, or apparatus that comprises such elements from including other same elements.

[00120] While the foregoing disclosure sets forth various embodiments using specific block diagrams, flowcharts, and examples, each block diagram component, flowchart step, operation, and/or component described and/or illustrated herein may be implemented, individually and/or collectively, using a wide range of hardware, software, or firmware (or any combination thereof) configurations. In addition, any disclosure of components contained within other components should be considered as examples because many other architectures can be implemented to achieve the same functionality.

[00121] The process parameters and sequence of steps described and/or illustrated herein are given by way of example only and can be varied as desired. For example, while the steps illustrated and/or described herein may be shown or discussed in a particular order, these steps do not necessarily need to be performed in the order illustrated or discussed. The various example methods described and/or illustrated herein may also omit one or more of the steps described or illustrated herein or include additional steps in addition to those disclosed.

[00122] While various embodiments have been described and/or illustrated herein in the context of fully functional computing systems, one or more of these example embodiments may be distributed as a program product in a variety of forms, regardless of the particular type of computer-readable medium used to actually carry out the distribution. The embodiments disclosed herein may also be implemented using software modules that perform certain tasks. These software modules may include script, batch, or other executable files that may be stored on a computer-readable storage media or in a computing system. These software modules may configure a computing system to perform one or more of the example embodiments disclosed herein. One or more of the software modules disclosed herein may be implemented in a cloud computing environment. Cloud computing environments may provide various services and applications via the Internet. These cloud-based services (e.g., software as a service, platform as a service, infrastructure as a service, etc.) may be accessible through a Web browser or other remote interface. Various functions described herein may be provided through a remote desktop environment or any other cloud-based computing environment.

[00123] Although the present disclosure and its advantages have been described in detail, it should be understood that various changes substitutions, and alterations can be made herein without departing from the spirit and scope of the disclosure as defined by the appended claims. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the disclosure and its practical applications, to thereby enable others skilled in the art to best utilize the disclosure and various embodiments with various modifications as may be suited to the particular use contemplated.

[00124] Moreover, the scope of the present application is not intended to be limited to the particular embodiments of the process, machine, manufacture, composition of matter, means, methods and steps described in the specification. As one of ordinary skill in the art will readily appreciate from the disclosure of the present disclosure, processes, machines, manufacture, compositions of matter, means, methods, or steps, presently existing or later to be developed, that

perform substantially the same function or achieve substantially the same result as the corresponding embodiments described herein may be utilized according to the present disclosure. Accordingly, the appended claims are intended to include within their scope such processes, machines, manufacture, compositions of matter, means, methods, or steps.

[00125] Embodiments according to the present disclosure are thus described. While the present disclosure has been described in particular embodiments, it should be appreciated that the disclosure should not be construed as limited by such embodiments, but rather construed according to the below claims.

WHAT IS CLAIMED IS:

1. A method of predicting GPU malfunctions, the method comprising:
installing a daemon program at a GPU node, the daemon program periodically collecting GPU status parameters corresponding to the GPU node at a pre-determined time period;
5 obtaining the GPU status parameters from the GPU node; and
comparing the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, wherein the mean status fault parameters are obtained by use of a pre-configured statistical model.
2. The method of claim 1, wherein the status parameter is temperature, and comparing the
10 obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction comprises:
performing statistics to obtain a temperature count of a number of times the GPU incurs a temperature greater than a pre-determined temperature threshold and a GPU temperature standard deviation threshold;
15 comparing the obtained temperature count with a mean temperature fault count;
comparing a GPU temperature fault standard deviation with the GPU temperature standard deviation threshold;
if the temperature count is greater than the mean temperature fault count, and if the GPU temperature fault standard deviation is less than the GPU temperature standard deviation threshold,
20 determining that the GPU is to malfunction; and
if the temperature count is less than the mean temperature fault count, or if the GPU temperature fault standard deviation is greater than the GPU temperature standard deviation threshold, determining that the GPU is not to malfunction.
3. The method of claim 2, wherein:
25 the GPU status parameters collected from the GPU node comprise a GPU model, temperature, and usage status; and
the obtaining mean status fault parameters by use of a pre-configured statistical model comprises:

determining, based on the usage status, whether the GPU malfunctions;

in response to a determination that the GPU does not malfunction, based upon the GPU model, storing the temperature collected from the GPU node in an information storage space corresponding to the GPU model; and

5 in response to a determination that the GPU malfunctions, based on the GPU model, obtaining temperatures stored in an information storage space corresponding to the GPU model, and based on the temperature collected from the GPU node and the stored temperatures obtained from the storage space corresponding to the GPU model, performing statistics to compute an arithmetic mean temperature fault count and a temperature fault
10 standard deviation, by use of a pre-configured temperature statistical model.

4. The method of claim 3, wherein the pre-configured temperature statistical model comprises:
a mean temperature fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model; and

15 a temperature fault standard deviation model configured to compute a temperature fault standard deviation based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, and the mean temperature fault count.

5. The method of claim 1, wherein the status parameter is power consumption, and comparing
20 the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction comprises:

performing statistics to obtain a power consumption count of a number of times the GPU incurs power consumption greater than a pre-determined power consumption threshold and a GPU power consumption standard deviation threshold;

25 comparing the obtained power consumption count with a mean power consumption fault count;

comparing a GPU power consumption fault standard deviation with the GPU power consumption standard deviation threshold;

if the power consumption count is greater than the mean power consumption fault count, and if the GPU power consumption fault standard deviation is less than the GPU power consumption standard deviation threshold, determining that the GPU is to malfunction; and

5 if the power consumption count is less than the mean power consumption fault count, or if the GPU power consumption fault standard deviation is greater than the GPU power consumption standard deviation threshold, determining that the GPU is not to malfunction.

6. The method of claim 5, wherein:

the GPU status parameters collected from the GPU node comprise a GPU model, power consumption, and usage status; and

10 the obtaining mean status fault parameters by use of a pre-configured statistical model comprises:

determining, based on the usage status, whether the GPU malfunctions;

15 in response to a determination that the GPU does not malfunction, based upon the GPU model, storing the power consumption collected from the GPU node in an information storage space corresponding to the GPU model; and

20 in response to a determination that the GPU malfunctions, based on the GPU model, obtaining power consumption stored in an information storage space corresponding to the GPU model, and based on the power consumption collected from the GPU node and the stored power consumption obtained from the storage space corresponding to the GPU model, performing statistics to compute an arithmetic mean power consumption fault count and a power consumption fault standard deviation, by use of a pre-configured power consumption statistical model.

7. The method of claim 6, wherein the pre-configured power consumption statistical model comprises:

25 a mean power consumption fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model; and

a power consumption fault standard deviation model configured to compute a power consumption standard deviation based on the GPU power consumption collected from the GPU

node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model, and the mean power consumption fault count.

8. The method of claim 1, wherein the status parameter is usage duration, and comparing the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction comprises:

comparing the obtained GPU usage duration with a mean fault usage duration;

if the GPU usage duration is greater than the mean fault usage duration, determining the GPU is to malfunction; and

if the GPU usage duration is less than the mean fault usage duration, determining the GPU is not to malfunction.

9. The method of claim 8, wherein:

the GPU status parameters collected from the GPU node comprise a GPU model, usage duration, and usage status; and

the obtaining mean status fault parameters by use of a pre-configured statistical model comprises:

determining, based on the usage status, whether the GPU malfunctions;

in response to a determination that the GPU does not malfunction, based upon the GPU model, storing the usage duration collected from the GPU node in an information storage space corresponding to the GPU model; and

in response to a determination that the GPU malfunctions, based on the GPU model, obtaining usage duration stored in an information storage space corresponding to the GPU model, and based on the usage duration collected from the GPU node and the stored usage duration obtained from the storage space corresponding to the GPU model, performing statistics to compute an arithmetic mean fault usage duration, by use of a pre-configured usage duration statistical model.

10. The method of claim 9, wherein the pre-configured usage duration statistical model comprises a mean fault usage duration model configured to compute an arithmetic mean based on the GPU usage duration collected from the GPU node and the stored GPU usage duration obtained from the information storage space corresponding to the GPU model.

11. An apparatus for predicting GPU malfunctions, the apparatus comprising:

a processor; and

a non-transitory computer-readable medium operably coupled to the processor, the non-transitory computer-readable medium having computer-readable instructions stored thereon to be

executed when accessed by the processor, the instructions comprising:

an installation module configured to install a daemon program at a GPU node, the daemon program periodically collecting GPU status parameters corresponding to the GPU node at a pre-determined time period;

a collecting module configured to obtain the GPU status parameters from the GPU node; and

a processing module configured to compare the obtained GPU status parameters with mean status fault parameters to determine whether the GPU is to malfunction, wherein the mean status fault parameters are obtained by use of a pre-configured statistical model.

12. The apparatus of claim 11, wherein the status parameter is temperature, and the processing module comprises:

a first statistical module configured to perform statistics to obtain a temperature count of a number of times the GPU incurs a temperature greater than a pre-determined temperature threshold and a GPU temperature standard deviation threshold;

a first comparison module configured to compare the obtained temperature count with a mean temperature fault count and to compare a GPU temperature fault standard deviation with the GPU temperature standard deviation threshold;

a first determination module configured to, if the temperature count is greater than the mean temperature fault count, and if the GPU temperature fault standard deviation is less than the GPU temperature standard deviation threshold, determine that the GPU is to malfunction; and

a second determination module configured to, if the temperature count is less than the mean temperature fault count, or if the GPU temperature fault standard deviation is greater than the GPU temperature standard deviation threshold, determine that the GPU is not to malfunction.

13. The apparatus of claim 12, wherein after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period; wherein the

collecting module comprises a first collecting module configured to collect from the GPU node a GPU model, temperature, and usage status; and wherein the processing module further comprises:

a first decision module configured to decide, based on the usage status, whether the GPU malfunctions;

5 a first storing module configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the temperature collected from the GPU node in an information storage space corresponding to the GPU model; and

a first computing module configured to, in response to a determination that the GPU malfunctions, based on the GPU model, obtain temperatures stored in an information storage space
10 corresponding to the GPU model, and based on the temperature collected from the GPU node and the stored temperatures obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean temperature fault count and a temperature fault standard deviation, by use of a pre-configured temperature statistical model.

14. The apparatus of claim 13, wherein the pre-configured temperature statistical model
15 comprises:

a mean temperature fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model; and

a temperature fault standard deviation model configured to compute a temperature fault
20 standard deviation based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model, and the mean temperature fault count.

15. The apparatus of claim 11, wherein the status parameter is power consumption, and wherein the processing module comprises:

25 a second statistical module configured to perform statistics to obtain a power consumption count of a number of times the GPU incurs power consumption greater than a pre-determined power consumption threshold and a GPU power consumption standard deviation threshold;

a second comparison module configured to compare the obtained power consumption count with a mean power consumption fault count and to compare a GPU power consumption fault
30 standard deviation with the GPU power consumption standard deviation threshold;

a third determination module configured to, if the power consumption count is greater than the mean power consumption fault count, and if the GPU power consumption fault standard deviation is less than the GPU power consumption standard deviation threshold, determine that the GPU is to malfunction; and

5 a fourth determination module configured to, if the power consumption count is less than the mean power consumption fault count, and or the GPU power consumption fault standard deviation is greater than the GPU power consumption standard deviation threshold, determine that the GPU is not to malfunction.

10 16. The apparatus of claim 15, wherein, after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period; wherein the collecting module comprises a second collecting module configured to collect from the GPU node a GPU model, power consumption, and usage status; and wherein the processing module further comprises:

15 a second decision module configured to, based on the usage status, decide whether the GPU malfunctions;

a second storing module configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the power consumption collected from the GPU node in an information storage space corresponding to the GPU model; and

20 a second computing module configured to, in response to a determination that the GPU malfunctions, based on the GPU model, obtain power consumption stored in an information storage space corresponding to the GPU model, and based on the power consumption collected from the GPU node and the stored power consumption obtained from the storage space corresponding to the GPU model, perform statistics to compute an arithmetic mean power consumption fault count and a
25 power consumption fault standard deviation, by use of a pre-configured power consumption statistical model.

17. The apparatus of claim 16, wherein the pre-configured power consumption statistical model comprises:

a mean power consumption fault count model configured to compute an arithmetic mean based on the GPU temperature collected from the GPU node and the stored GPU temperatures obtained from the information storage space corresponding to the GPU model; and

5 a power consumption fault standard deviation model configured to compute a power consumption fault standard deviation based on the GPU power consumption collected from the GPU node and the stored GPU power consumption obtained from the information storage space corresponding to the GPU model, and the mean power consumption fault count.

18. The apparatus of claim 11, wherein the status parameter is usage duration; and wherein the processing module comprises:

10 a third comparison module configured to compare the obtained GPU usage duration with a mean fault usage duration;

a fifth determination module configured to, if the GPU usage duration is greater than the mean fault usage duration, determine the GPU is to malfunction; and

15 a sixth determination module configured to, if the GPU usage duration is less than the mean fault usage duration, determine the GPU is not to malfunction.

19. The apparatus of claim 18, wherein, after the installation module installs the daemon programs on the GPU node, the daemon program periodically collects the GPU model and GPU status parameters corresponding to the GPU node at a pre-determined time period; wherein the collecting module comprises a third collecting module configured to collect from the GPU node a GPU model, usage duration, and usage status; and wherein the processing module further comprises:

20 a third decision module configured to, based on the usage status, decide whether the GPU malfunctions;

25 a third storing module configured to, in response to a determination that the GPU does not malfunction, based upon the GPU model, store the usage duration collected from the GPU node in an information storage space corresponding to the GPU model; and

a third computing module configured to, in response to a determination that the GPU malfunctions, based on the GPU model, obtain usage duration stored in an information storage space corresponding to the GPU model, and based on the usage duration collected from the GPU node and the stored usage duration obtained from the storage space corresponding to the GPU model, perform

statistics to compute an arithmetic mean fault usage duration, by use of a pre-configured usage duration statistical model.

20. The apparatus of claim 19, wherein the pre-configured usage duration statistical model comprises a mean fault usage duration model configured to compute an arithmetic mean based on the GPU usage duration collected from the GPU node and the stored GPU usage duration obtained from the information storage space corresponding to the GPU model.

1/8




FIG. 1

2/8




FIG. 2

3/8





FIG. 3

4/8



5/8




FIG. 5

6/8




FIG. 6

7/8




FIG. 7

8/8




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/019765

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 11/07 (2016.01)

CPC - G06F 11/3058 (2016.01)

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G05D 23/00; G06F 3/038, 9/44, 11/00, 07, 15/00, 17/30, 19/00; H02H 3/05; H04L 12/26 (2016.01)

CPC - G06F 1/3218, 3/14, 11/008, 073, 0754, 2257, 3058, 15/00, 17/30144; G09G 2330/02; H04L 43/16 (2016.01)

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

U.S. Class/Subclass(es): 345/212, 501; 702/34; 703/21; 707/694; 709/224; 714/4, 47 (keyword delimited)

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

Orbit, Google Patents, Google Scholar

Search terms used: GPU, graphics processor, malfunction, failure, predicting, daemon, status parameter, statistical model, temperature threshold, usage status, power consumption, usage duration

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014/0092103 A1 (NVIDIA CORPORATION) 03 April 2014 (03.04.2014) entire document	1-20
Y	US 2009/0106571 A1 (LOW et al) 23 April 2009 (23.04.2009) entire document	1-20
Y	US 6,286,992 B1 (KYRTSOS) 11 September 2001 (11.09.2001) entire document	2-7, 12-17
Y	US 6,308,138 B1 (JONES et al) 23 October 2001 (23.10.2001) entire document	5-7, 15-17
A	US 2004/0225911 A1 (SMITH) 11 November 2004 (11.11.2004) entire document	1-20
A	US 2010/0118019 A1 (CASES et al.) 13 May 2010 (13.05.2010) entire document	1-20
A	US 2010/0306598 A1 (ACKARET et al.) 02 December 2010 (02.12.2010) entire document	1-20
A	US 2011/0246093 A1 (WOOD et al.) 06 October 2011 (06.10.2011) entire document	1-20
A	US 2013/0138419 A1 (LOPEZ et al.) 30 May 2013 (30.05.2013) entire document	1-20
A	US 2014/0067779 A1 (OJHA) 06 March 2014 (06.03.2014) entire document	1-20
A	US 2013/0275582 A1 (CITRIX SYSTEM, INC.) 17 October 2013 (17.10.2013) entire document	1-20
A	US 2011/0313693 A1 (INOUE et al) 22 December 2011 (22.12.2011) entire document	1-20

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

11 April 2016

Date of mailing of the international search report

05 MAY 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774