



- (51) International Patent Classification: **G06F 19/00** (2011.01)
- (21) International Application Number: PCT/US2014/032424
- (22) International Filing Date: 31 March 2014 (31.03.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
 

|            |                              |    |
|------------|------------------------------|----|
| 61/807,665 | 2 April 2013 (02.04.2013)    | US |
| 14/063,982 | 25 October 2013 (25.10.2013) | US |
- (71) Applicant: **WESTERN DIGITAL TECHNOLOGIES, INC.** [US/US]; 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US).
- (72) Inventor: **AHMADI-ARDAKANI, Arya**; 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US).
- (74) Agents: **PUA, Don** et al.; Western Digital Technologies, Inc., 3355 Michelson Drive, Suite 100, Irvine, California 92612 (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, LT, 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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

[Continued on next page]

(54) Title: DETECTION OF USER BEHAVIOR USING TIME SERIES MODELING




FIG. 6

(57) Abstract: The embodiments provide a way to predict when a storage device will be accessed. In order to enhance performance, the storage device may proactively prepare for the access operation, and thus, minimize the access-time response of the storage device. The user behavior is recorded over time and collected into a dataset. In one embodiment, the intervals between the data points in the dataset are calculated and arranged into a matrix. Patterns in the matrix are recognized and used to recognize the next likely access by the user. The storage device may then take various actions, such as drive spin up, in anticipation of the next predicted access to minimize access-time response.

WO 2014/165456 A1



EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,  
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,  
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,  
GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

## DETECTION OF USER BEHAVIOR USING TIME SERIES MODELING

### BACKGROUND

**[0001]** The responsiveness of a computing device is an important characteristic of its performance. Ideally, a computing device is able to respond to a user as quickly as possible.

**[0002]** For example, storage devices attempt to provide access to their stored data as efficiently and as quickly as possible. Today, solid-state drives (SSDs) have become popular due to their near-zero data access-time delay (DATD) performance relative to hard disk drives (HDDs). HDDs, however, still offer higher capacity storage at a lower cost relative to SSDs. Accordingly, HDDs remain popular due to their high capacity and durability. Indeed, many known products, such as solid state hybrid drives (SSHD) now provide a combination of SSD and HDD as a storage solution.

**[0003]** HDDs normally have a high DATD because they are frequently put into a standby mode in order to conserve power, etc. For example, mobile devices, such as a laptop, wireless Network Attached Storage (NAS), or any other mobile device, frequently employ aggressive HDD standby policies. HDDs may also be put into a standby mode for data safety reasons during various movements the head would not make contact with the disk.

**[0004]** Unfortunately, when exiting standby mode in response to a data access request, an HDD may take a long time for the drive to spin back up. Moreover, as the hard drive capacities increase, the DATD may also increase when coming out of standby mode.

**[0005]** Therefore, it would be desirable to provide methods and systems that can minimize the DATD of a HDD. It would also be desirable to provide methods and systems that can improve the responsiveness of any computing device, especially after idle periods of time.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[0006]** Systems and methods which embody the various features of the invention will now be described with reference to the following drawings, in which:

**[0007]** FIG. 1 illustrates an exemplary system in which embodiments of the present invention may be implemented;

**[0008]** FIG. 1A illustrates an exemplary controller that is configured to predict user data requests in accordance with embodiments of the present invention;

**[0009]** FIG. 2 illustrates an exemplary user behavioral pattern;

**[0010]** FIG. 3 illustrates an exemplary predicted access event based on time series modeling by some embodiments of the present invention;

**[0011]** FIG. 4 illustrates exemplary time series modeling used by some embodiments to predict various data access events;

**[0012]** FIG. 5 illustrates an exemplary graph of data access events predicted by the some embodiments and actual user data requests; and

**[0013]** FIG. 6 illustrates an exemplary process flow.

## **DETAILED DESCRIPTION**

### **[0014] Overview – Predicting User Behavior and Data Access Requests**

**[0015]** Embodiments of the present invention improve the data access time delay for accessing stored data based on predicting when data access requests may occur. In some embodiments, the storage device observes user behavior and uses time series modeling to predict or anticipate when a user is likely to access a storage system. Some embodiments are implemented on storage devices containing one or more hard disk drives alone or in combination with other types of storage devices. The embodiments may employ software, firmware, etc. to implement the predictive algorithms used to anticipate when data access requests may occur.

**[0016]** For the purpose of illustration, an example will now be described with reference to a network-attached storage that comprises one or more hard disk drives. Those skilled in the art will recognize that the embodiments may be implemented on a wide variety of types of storage types, such as direct-attached storage, storage area networks, and the like.

**[0017]** Certain embodiments will now be described. These embodiments are presented by way of example only, and are not intended to limit the scope of the inventions. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions, substitutions

and changes in the form of the methods and systems described herein may be made without departing from the spirit of the inventions. To illustrate some of the embodiments, reference will now be made to the figures.

**[0018]** FIG. 1 illustrates an exemplary system. FIG. 1A illustrates an exemplary controller that is configured to predict user data requests in accordance with embodiments of the present invention. FIG. 2 illustrates an exemplary user behavioral pattern, which may be analyzed and predicted by the embodiments. FIG. 3 illustrates an exemplary predicted access event based on time series modeling by some embodiments of the present invention. FIG. 4 illustrates exemplary time series modeling used by some embodiments to predict various data access events. FIG. 5 illustrates an exemplary graph of data access events predicted by some embodiments and actual user data requests. FIG. 6 illustrates an exemplary process flow. FIGs. 1-5 will now be further described below.

**[0019] Exemplary System – FIG. 1**

**[0020]** FIG. 1 shows an exemplary system of an embodiment of the present invention. The embodiments of the present invention may be implemented on any type of storage device or device that provides file notification services. For purposes of illustration, an exemplary system 100 is shown with a network-attached storage (NAS) 102 that services a client (or host device) 104 via a network 106. The system 100 may also be used in other devices in which predictive capabilities are desired, such as mobile devices, media streaming/playback devices, laptops, tablets, desktop computers, smart televisions, or other types of electronic devices. An exemplary embodiment with these components will now be briefly described below.

**[0021] Exemplary Storage Device – NAS 102**

**[0022]** The NAS 102 is file-level storage device that is capable of connecting to a network, such as network 106 and provides access to files stored in its storage medium as a file server. In one embodiment, NAS 102 is implemented with known hardware, software, and firmware. For example, in one embodiment, the NAS 102 is configured as an appliance having an embedded operating system. The NAS 102 may support a variety of operating systems, such as UNIX, LINUX, Windows, and the like. As will be further described, the NAS 102 may also comprise multiple storage mediums, such as one or more hard disks that are arranged into a

redundant array of independent disks (RAID). Furthermore, the NAS 102 may support various protocols, such as NFS, SMB/CIFS, AFP, etc.

**[0023]** As shown, the NAS 102 may comprise a network interface 110, a controller 112, a storage interface 114, and a storage array of disks 116. These components will now be briefly described below.

**[0024]** Network interface 110 serves as the network communications interface for the NAS 102 to/from network 106. For example, in one embodiment, the network interface 200 may comprise one or more Gigabit Ethernet, Ethernet, Universal Serial Bus (USB), or other interfaces for communications with network 106. Such components are known to those skilled in the art.

**[0025] Controller 112**

**[0026]** Controller 112 represents the hardware and software that manages the disks 116 of the NAS 102 and presents them as a logical unit or volume to the clients 104. In some embodiments, the controller 112 may also comprise one or more other components to supplement its operations, such as an on-chip RAID controller, a memory or disk cache, etc. Furthermore, in the embodiments, the controller 112 is configured to observe user behavior with regard to data requests from client 104, detect patterns in that behavior, and predict various future data requests. The controller 112 may then manage or manipulate the operations and power state of the various disks 116 via the storage interface 114. In some embodiments, the controller 112 employs a time series modeling to predict future data requests and manages the disks 116 accordingly to improve their data access time and/or minimize their DATD. An example of the time series modeling algorithms and implementations is described further below.

**[0027] Storage Interface 114**

**[0028]** Storage interface 114 serves as an interface between the controller 112 and the disks 206. The storage interface 114 may support various communications, such as SAS, SATA, SCSI, etc.

**[0029] Disk 116**

**[0030]** Disks 116 represent the storage medium and associated electronics for the devices storing data for the NAS 102. In one embodiment, the disks 116 may

be implemented as hard disk drives, such as those provided by Western Digital Technologies, Inc. Of course, in other embodiments, the NAS 102 may comprise other types of storage media and devices, such as solid-state drives, hybrid drives, etc. Any type of storage drive that can be configured as part of a RAID may be implemented as part of an embodiment of the present invention.

**[0031]** Alternatively, as noted above, in other embodiments, the system 100 may comprise direct attached storage devices or external drives. For example, the system 100 may comprise one or more external drives, such as the MyBook or Passport devices provided by Western Digital Technologies, Inc.

**[0032] Clients 104**

**[0033]** Clients 104 represent the various client devices that may store and/or access files on the NAS 102. For example, the clients 104 may be a desktop, a laptop, a tablet, a smart phone, etc. The embodiments support any device that can access a file stored on the NAS 102.

**[0034] Network 106**

**[0035]** Network 106 provides a communication infrastructure for data communications between the components of system 100. Network 106 may comprise known network elements, such as hubs, switches, routers, firewalls, etc., to facilitate and secure these communications. In the embodiments, the network 106 may comprise a local area network, a wide area network, etc. In addition, the network 106 may comprise wired and wireless links or components to carry its communications.

**[0036] Exemplary Controller Implementing Time Series Modeling – FIG.**

**1A**

**[0037]** Further to FIG. 1, FIG. 1A shows a more detailed block diagram of the controller 112 that is configured to predict future data requests based on time series modeling. In some embodiments, the controller 112 may implement a time series modeler that comprises four main components: (1) an event recorder 118; (2) a detection algorithm 120; (3) a predictor algorithm 122; and (4) a power manager 124. These components will now be further described.

**[0038] Event Recorder 118**

**[0039]** The event recorder 118 records information about the data requests received from the client 104. The event recorder 118 may record various aspects of the data requests, such as date/time of the request, file type, file size, etc. The event recorder 118 may be implemented in the controller 112 in a wide variety of ways. For example, the event recorder 118 may be firmware within the controller 112 that is configured to perform the functions of recording data requests. As another example, the event recorder 118 may be software running in the controller 112 that continually records the behavior of the user of the device over time. In some embodiments, the event recorder 118 records some or all of the data requests from the clients 104 and every data point (or a selection of data points) essentially serves as a training point and the collection of data points continually adds to the accuracy of the overall prediction.

**[0040] Detection Algorithm 120**

**[0041]** The controller 112 executes the detection algorithm 120 to recognize various patterns within the previously data collected by the event recorder 118. In some embodiments, the detection algorithm 120 determines the various inter-arrival intervals of the data requests to generate a set of time series of intervals. The detection algorithm 120 may then identify or detect user behavior based on the patterns existing in these time series of intervals.

**[0042] Predictor Algorithm 122**

**[0043]** The controller 112 may also execute a predictor algorithm 122. Once a pattern has emerged, the predictor algorithm 122 provides the logic for predicting the future behavior of the user and when one or more future data requests from client 104 is expected. The predictor algorithm 122 may output its predictions in various formats, such as a time offset, a confidence factor, etc.

**[0044] Power Manager 124**

**[0045]** Finally, the controller 112 may comprise a power manager 124. The power manager 124 manages the power state and power consumption of the NAS 102, such as the storage interface 114 and disks 116. In some embodiments, the controller 112 is configured to use the power manager 124 to anticipate user data request usage based on the predictions of the predictor algorithm 122. For example,

in response to a predicted future data request by the predictor algorithm 122, the controller 112 may signal the power manager 124 to transition the power state of the disks 116 to commence spinning the hard disks 116 at the predicted time of the next data request. In addition, the controller 112 may signal the power manager 124 to power down or reduce the power consumption of the disks 116 and/or the storage interface 114 when a long idle period or time until the next data request is predicted. Various thresholds for when the controller 112 and the power manager 124 change the power state and operation may vary, for example, based on spin up times, communications latency, read/write latency, etc. Furthermore, the controller 112 and the power manager 124 may modify its power management schemes based on various criteria, such as battery charge state, frequency of data requests, and the like.

**[0046] Example of Time Series Modeling**

**[0047]** To illustrate the concepts of the invention, the detection and prediction algorithms will now be described further below with reference to FIGs 2-5. For the purpose of illustration, a simple example of a user behavior and series of data requests is shown with reference to FIG. 2.

**[0048] FIG. 2 – Example User Behavior (Simplified)**

**[0049]** As shown in FIG. 2, the graph can represent the behavior of a user at client 104 that has accessed data on the device in 100 hour intervals since he has turned on the NAS 102 initially. In this example, the controller 112 can thus make a simple prediction can from such user behavior.

**[0050] FIG. 3 – Prediction of User Behavior**

**[0051]** For example, FIG. 3 illustrates a predicted data request event determined by the controller 112. As shown, various “User-Data-Request” points (to indicate data requests from the client 104) are marked as  $R_n$  and a “Future Event Prediction” point predicted by the controller 112 is marked as  $P_n$  where  $n$  is 1,2,3, and so forth. As can be seen, the time interval between  $R_2$  and  $R_1$  may be represented as  $\Delta_1$  and  $R_3$  and  $R_2$  may be represented as  $\Delta_2$ .

**[0052]** In one embodiment, if  $\Delta_1$  and  $\Delta_2$  are found to be equal or within a certain tolerance, the predictor algorithm 122 can predict a pattern and hence a

future data request. That is, the next user request can be predicted  $\Delta_1$  or  $\Delta_2$  away from  $R_3$  (i.e. about 100 hours).

**[0053] Accounting for Fluctuations and/or Variations**

**[0054]** In some embodiments, in order to account for some fluctuations in user request behavior, the prediction algorithm 122 may include an error buffer. For example, the algorithm may be formulated and simplified below:

if  $|\Delta_{n+1} - \Delta_n| \leq \epsilon_r \Rightarrow \Delta'_{n+2} = \min(\Delta_n, \Delta_{n+1}) - \epsilon_r - t_s$ , where *min* is the minimum function,

$\epsilon_r$  is the buffer accounting for fluctuations in user request, and

$t_s$  is the DATD, such as the time it takes for the drive to spin up and device to get ready to respond to user's request), and finally

$\Delta'_{n+2}$  is the future prediction interval (shown in FIG. 3 as the interval between  $R_3$  and  $P_1$ ).

**[0055]** In one embodiment, for the ease of computation, the detection algorithm 120 is configured to express the deltas in matrix form as suggested below:

$$\begin{bmatrix} \Delta_{1,1} & \cdots & \Delta_{1,j} \\ \vdots & \ddots & \vdots \\ \Delta_{i,1} & \cdots & \Delta_{i,j} \end{bmatrix}$$

where  $\Delta_{i,j}$  refers to the time interval between  $R_i$  and  $R_j$ .

**[0056]** As can be seen, the matrix is a square matrix with all the elements on its main diagonal equaling 0. In addition, the lower triangular part of the matrix contains the same data as the upper triangular part of the matrix so that lower triangular part can be neglected (for example to conserve memory) and based on these the simplification of the above matrix the following matrix results:

$$\begin{bmatrix} 0 & \Delta_{1,2} & \cdots & \Delta_{1,j} \\ 0 & 0 & \cdots & \Delta_{2,j} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}$$

**[0057]** Accordingly, any row  $i$  contains all deltas pertaining to a data request point  $R_i$  to all other future data request points. Also any column  $j$  contains all deltas pertaining to intervals from all other past data request point to  $R_j$ . In order to visualize this example, FIG. 4 shows the graph illustrating the intervals expressed in the matrix.

**[0058] Time Series Modeling – Additional Example**

**[0059]** In some embodiments, the controller 112 is configured to perform time series modeling to identify patterns within the matrix, based on the prediction algorithm 122. An example of the detection algorithm 120 and time series modeling will now be provided in reference to the following hypothetical user pattern.

| User Data Request points | R <sub>1</sub> | R <sub>2</sub> | R <sub>3</sub> | R <sub>4</sub> | R <sub>5</sub> | R <sub>6</sub> | R <sub>7</sub> | R <sub>8</sub> | R <sub>9</sub> | R <sub>10</sub> | R <sub>11</sub> | R <sub>12</sub> |
|--------------------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|-----------------|-----------------|-----------------|
| Unit of Time (hrs)       | 5              | 7              | 10             | 19             | 21             | 24             | 33             | 35             | 38             | 48              | 58              | 68              |

**[0060]** From this dataset, the controller 112 may construct the following simplified matrix below:

|   |   |   |    |    |    |    |    |    |    |    |    |
|---|---|---|----|----|----|----|----|----|----|----|----|
| 0 | 2 | 5 | 14 | 16 | 19 | 28 | 30 | 33 | 43 | 53 | 63 |
| 0 | 0 | 3 | 12 | 14 | 17 | 26 | 28 | 31 | 41 | 51 | 61 |
| 0 | 0 | 0 | 9  | 11 | 14 | 23 | 25 | 28 | 38 | 48 | 58 |
| 0 | 0 | 0 | 0  | 2  | 5  | 14 | 16 | 19 | 29 | 39 | 49 |
| 0 | 0 | 0 | 0  | 0  | 3  | 12 | 14 | 17 | 27 | 37 | 47 |
| 0 | 0 | 0 | 0  | 0  | 0  | 9  | 11 | 14 | 24 | 34 | 44 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 2  | 5  | 15 | 25 | 35 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 3  | 13 | 23 | 33 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 10 | 20 | 30 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 10 | 20 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 10 |

**[0061]** In one embodiment, the detection algorithm 120 may be implemented based on the following pseudo-code to identify patterns in the user access request.

```

clc
G=[];
I=[5 7 10 19 21 24 33 35 38 48 58 68]; % vector of intervals starting at
i=1

I_size=size(I);
for i=1:I_size;
for j=1:I_size

```

```

Delta(i,j)=abs(l(i)-l(j));
end
end
Delta=triu(Delta)

for i=1:l_size
col_v=Delta(:,i);
row_v=Delta(i,:);

t_col=col_v(col_v~=0);
t_row=row_v(row_v~=0);

RESULT=ismember(col_v,t_row);
[row_i,col_i]=find(RESULT == 1);
if isempty(Delta(row_i,i))
% disp 'this is empty'
% Delta(row_i,i)
else
detected_delta=Delta(row_i,i)
G(end+1)=l(i)+detected_delta;
end
end
figure
plot(l,1,'bx', G, 1,'ro');
legend('User Data Request','Future Event Prediction');

```

**[0062]** The above code also performs the prediction of the user behavior and produces the following graph as shown in FIG. 5. As can be seen, the user has interacted with the NAS 102 with data requests that arrive in a distinct pattern. The exemplary matrix below shows a pattern that has been identified by the detection algorithm 120 matching series of intervals that are common to both a row and a column from a selected element of the matrix.

|   |   |   |    |    |    |    |    |    |    |    |    |
|---|---|---|----|----|----|----|----|----|----|----|----|
| 0 | 2 | 5 | 14 | 16 | 19 | 28 | 30 | 33 | 43 | 53 | 63 |
| 0 | 0 | 3 | 12 | 14 | 17 | 26 | 28 | 31 | 41 | 51 | 61 |
| 0 | 0 | 0 | 9  | 11 | 14 | 23 | 25 | 28 | 38 | 48 | 58 |
| 0 | 0 | 0 | 0  | 2  | 5  | 14 | 16 | 19 | 29 | 39 | 49 |
| 0 | 0 | 0 | 0  | 0  | 3  | 12 | 14 | 17 | 27 | 37 | 47 |
| 0 | 0 | 0 | 0  | 0  | 0  | 9  | 11 | 14 | 24 | 34 | 44 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 2  | 5  | 15 | 25 | 35 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 3  | 13 | 23 | 33 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 10 | 20 | 30 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 10 | 20 |
| 0 | 0 | 0 | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 10 |

**[0063]** In some embodiments, some or all of above system and its components are executed or embodied by one or more physical or virtual computing systems. For example, in some embodiments, a computing system (such as a server, mobile device or other computing device) that has components including a central processing unit (CPU), input/output (I/O) components, storage and/or memory may be used to execute some or all of the components of the above system. The I/O components can include a display, a network connection to the network, a computer-readable media drive and/or other I/O devices (e.g., a keyboard, a mouse, a touch screen, speakers, etc.).

**[0064]** FIG. 6 illustrates an exemplary process flow for predicting future access of a storage device, such as storage device 102. In stage 600, the event recorder 118 in storage device 102 records data requests are received for files stored on the disks 116. The controller 112 then determines respective intervals between each access request to any subsequent access request of the storage device 102. For example, the controller 112 may generate a data structure, such as an array or matrix, wherein each element indicates the interval (or delta). In some embodiments, the matrix may be further simplified or reduced because the matrix is a square matrix, and thus, the lower triangular half can be neglected. As noted, the matrix may be stored in a memory (not shown) by controller 112. An example of such a matrix is shown above.

**[0065]** In stage 602, the controller 112 identifies a pattern in the respective intervals. For example, as noted above, the controller 112 may execute program code or coordinate with a hardware or firmware component that implements a prediction algorithm, such as predictor algorithm 122 described above.

**[0066]** In stage 604, the controller 112 predicts a next access request based on the pattern in the respective intervals. In particular, the controller 112 determines a time for the next data request, based on the output of the predictor algorithm. The predicted next data request may be provided as an absolute time, a time offset, a range, etc. The predicted next data request may also have associated meta-data, such as a confidence factor, the type of data request, etc.

**[0067]** In stage 604, the controller 112 may issue one or more commands based on the next anticipated data request. For example, the controller 112 may activate or deactivate the proactively activating the storage device based on the predicted next access request. For example, the controller 112 may transition the disks 116 from one power state to another power state, such as active to standby. The controller 112 may also activate or deactivate other components, such as the network interface 110, the storage interface 114, etc., for example, to conserve battery consumption or maximize performance. Furthermore, the controller 112 may command the caching or pre-fetching of various files stored on the disks 116 based on the next predicted data request.

**[0068]** In one embodiment, the controller 112 may be configured to employ the time series modeling to optimize battery life of the storage device 102. For example, the controller 112 may comprise a power manager 124 that controls the power management of the storage device 102, especially the power consumption of the disks 116. In some embodiments, the power manager 124 may transition the power state of disks 116 to commence spinning down or up of the disks 116 when a long idle time is requested or when a next data request is predicted.

**[0069]** Furthermore, the controller 112 may also account for various tolerances, such as spin up delay of the disks 116, etc. The tolerances may be based on known characteristics of the disks 116 or based on meta-data provided by the detection algorithm 120.

**[0070]** Each of the processes, methods and algorithms described in the preceding sections may be embodied in, and fully or partially automated by, code modules executed by one or more computers, computer processors, or machines configured to execute computer instructions. The code modules may be stored on any type of non-transitory computer-readable storage medium or tangible computer storage device, such as hard drives, solid state memory, optical disc and/or the like. The processes and algorithms may be implemented partially or wholly in application-specific circuitry. The results of the disclosed processes and process steps may be stored, persistently or otherwise, in any type of non-transitory computer storage such as, e.g., volatile or non-volatile storage.

**[0071]** The features and attributes of the specific embodiments disclosed above may be combined in different ways to form additional embodiments, all of which fall within the scope of the present disclosure. Although the present disclosure provides certain embodiments and applications, other embodiments that are apparent to those of ordinary skill in the art, including embodiments, which do not provide all of the features and advantages set forth herein, are also within the scope of this disclosure. Accordingly, the scope of the present disclosure is intended to be defined only by reference to the appended claims.

**WHAT IS CLAIMED IS:**

1. A method of predicting future access of a storage device, said method comprising:

determining respective intervals between each access request to any subsequent access request of the storage device;

identifying a pattern in the respective intervals;

predicting a next access request based on the pattern in the respective intervals; and

proactively activating the storage device based on the predicted next access request.

2. The method of claim 1, wherein proactively activating the storage device comprises proactively spinning up a hard disk drive before the predicted next access request based on a spin up time of the hard disk.

3. The method of claim 1, wherein proactively activating the storage device comprises proactively activating the storage device before the predicted next access request based on the respective intervals and a tolerance.

4. The method of claim 1, further comprising determining an error tolerance between the respective intervals.

5. The method of claim 4, further comprising grouping a plurality of access requests that occur within the error tolerance.

6. The method of claim 1, further comprising arranging the respective intervals in a matrix.

7. The method of claim 6, wherein identifying the pattern in the respective intervals determining a first set of intervals occurring in a column of the matrix that correspond to a second set of intervals occurring in a row of the matrix.

8. A method of managing an electronic device based on predicting user behavior, said method comprising:

recording a series of user access requests of the electronic device over a period of time;

determining respective intervals between the series of user access requests;

identifying a pattern of user access requests based on modeling the series of user access requests based on matching a series of intervals of the user access requests;

determining a predicted time of a next user access request based on the identified pattern; and

initiating an action on the electronic device based on the predicted time of the next user access request.

9. The method of claim 8, wherein identifying the pattern of user access requests comprises:

arranging the respective intervals into a matrix wherein each column of the matrix indicates intervals from all other past user access requests and wherein each row indicates intervals to subsequent user access requests; and

identifying the pattern of the user access requests based on identifying a sequence of intervals that occur in both a row and column from a selected element in the matrix.

10. The method of claim 9, wherein identifying the pattern of user access requests comprises identifying a sequence of intervals that occur in both a row and column from a selected element in the matrix within a tolerance.

11. The method of claim 8, wherein initiating the action on the electronic device comprises initiating spin up of a hard disk drive.

12. The method of claim 8, wherein initiating the action on the electronic device comprises initiating a spin down of a hard disk drive.

13. The method of claim 8, wherein initiating the action on the electronic device comprises transitioning a storage drive from a standby state to an active state.

14. The method of claim 8, wherein initiating the action on the electronic device comprises transitioning a storage drive from an active state to a standby state.

15. A storage device configured to anticipate user behavior based on modeling a time series of user access requests, said storage device comprising:

at least one storage medium;

a controller comprising:

an event recorder configured to record a time series of user access requests for data from the at least one storage medium;

a power manager configured to manage power consumption of the at least one storage medium based on an active state and a standby state;

a time series modeler configured to detect patterns of user behavior based on intervals in the time series between user access requests; and

a predictor configured to predict a next user access request based on the patterns of user behavior and send a signal to the power manager to transition the storage device from a first power state to a second power state in anticipation of the predicted next user access request.

16. The storage device of claim 15, further comprising a memory configured to store a matrix that indicates intervals between user access requests.

17. The storage device of claim 16, wherein the time series modeler is configured to detect patterns of user behavior based on matching sequences of intervals occurring in both a row and a column of the matrix.

18. The storage device of claim 15, wherein predictor is configured to send a signal to transition the at least on storage medium from a standby state to an active state in anticipation of the predicted next user access request.

19. The storage device of claim 15, wherein predictor is configured to send a signal to transition the at least on storage medium from an active state to a standby state in anticipation of the predicted next user access request.

20. A device configured to predict user behavior, said device comprising:  
a storage storing information subject to user access requests; and  
a processor configured by executable program code to record a series of user access requests of the device over a period of time, determine respective intervals between the series of user access requests, identify a pattern of user access requests based on modeling the series of user access requests based on matching a series of intervals of the user access requests, determine a predicted time of a next user access request based on the identified pattern, and initiate an action on the device based on the predicted time of the next user access request.

21. The device of claim 20, wherein the processor is configured to arrange the respective intervals into a matrix stored in a memory coupled to the processor, wherein each column of the matrix indicates intervals from all other past user access requests and wherein each row indicates intervals to subsequent user access requests; and

wherein the processor is configured to identify the pattern of the user access requests based on identifying a sequence of intervals that occur in both a row and column from a selected element in the matrix.

22. The device of claim 20, wherein the storage comprises a hard disk drive and the processor is configured to initiate spin up of a hard disk drive.

23. The device of claim 20, wherein the storage comprises a hard disk drive and the processor is configured to initiate spin down of a hard disk drive.

24. The device of claim 20, wherein the storage comprises a hard disk drive and the processor is configured to transition the hard disk drive from a standby state to an active state.

25. The device of claim 20, wherein the storage comprises a hard disk drive and the processor is configured to transition the hard disk drive from an active state to a standby state.

26. The device of claim 20, wherein the processor is configured to initiate caching content responsive to the pattern of user access.

27. The device of claim 20, wherein the processor is configured to connect to a network responsive to the pattern of user access.

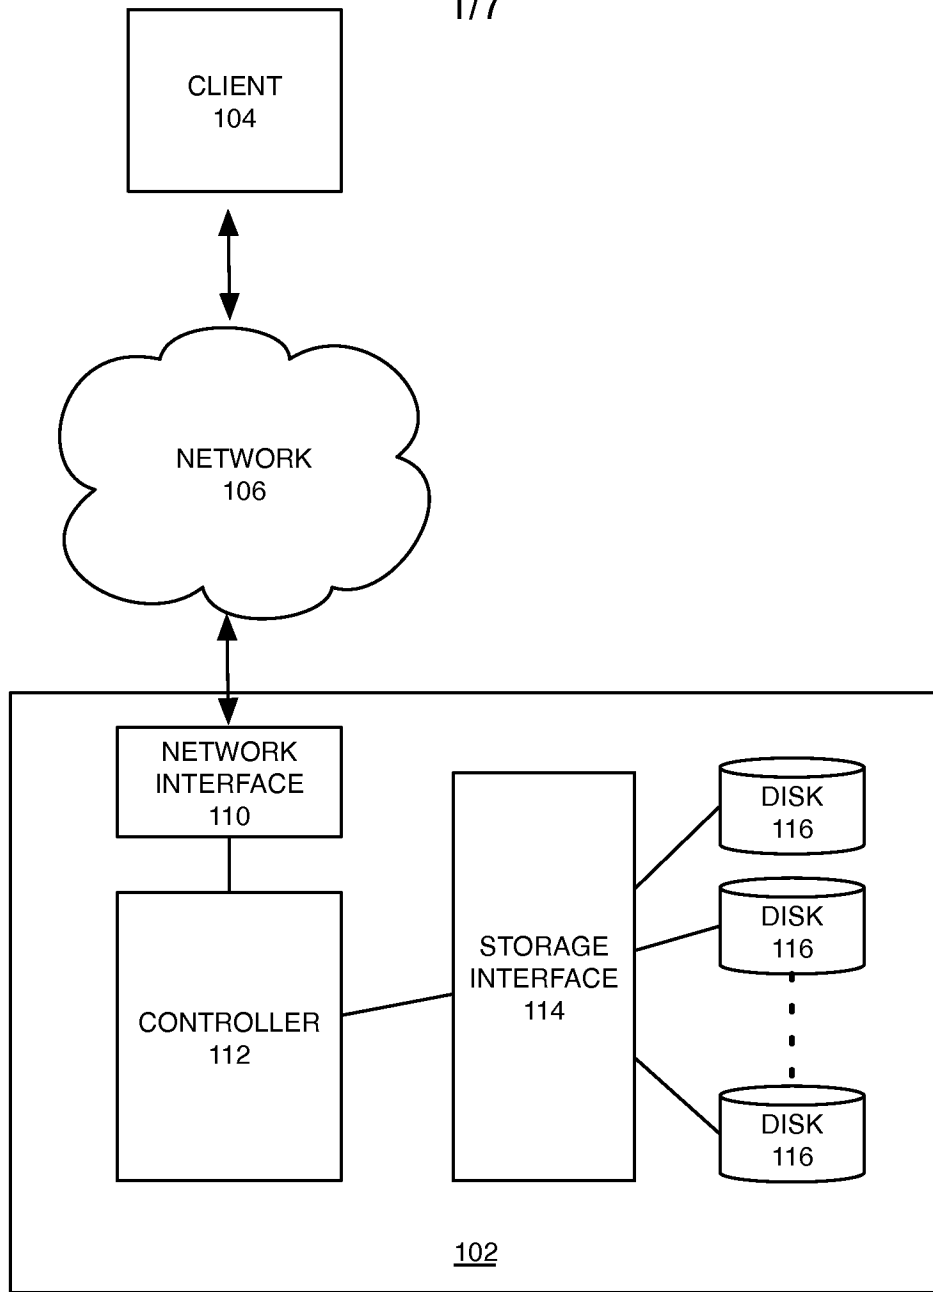


FIG. 1

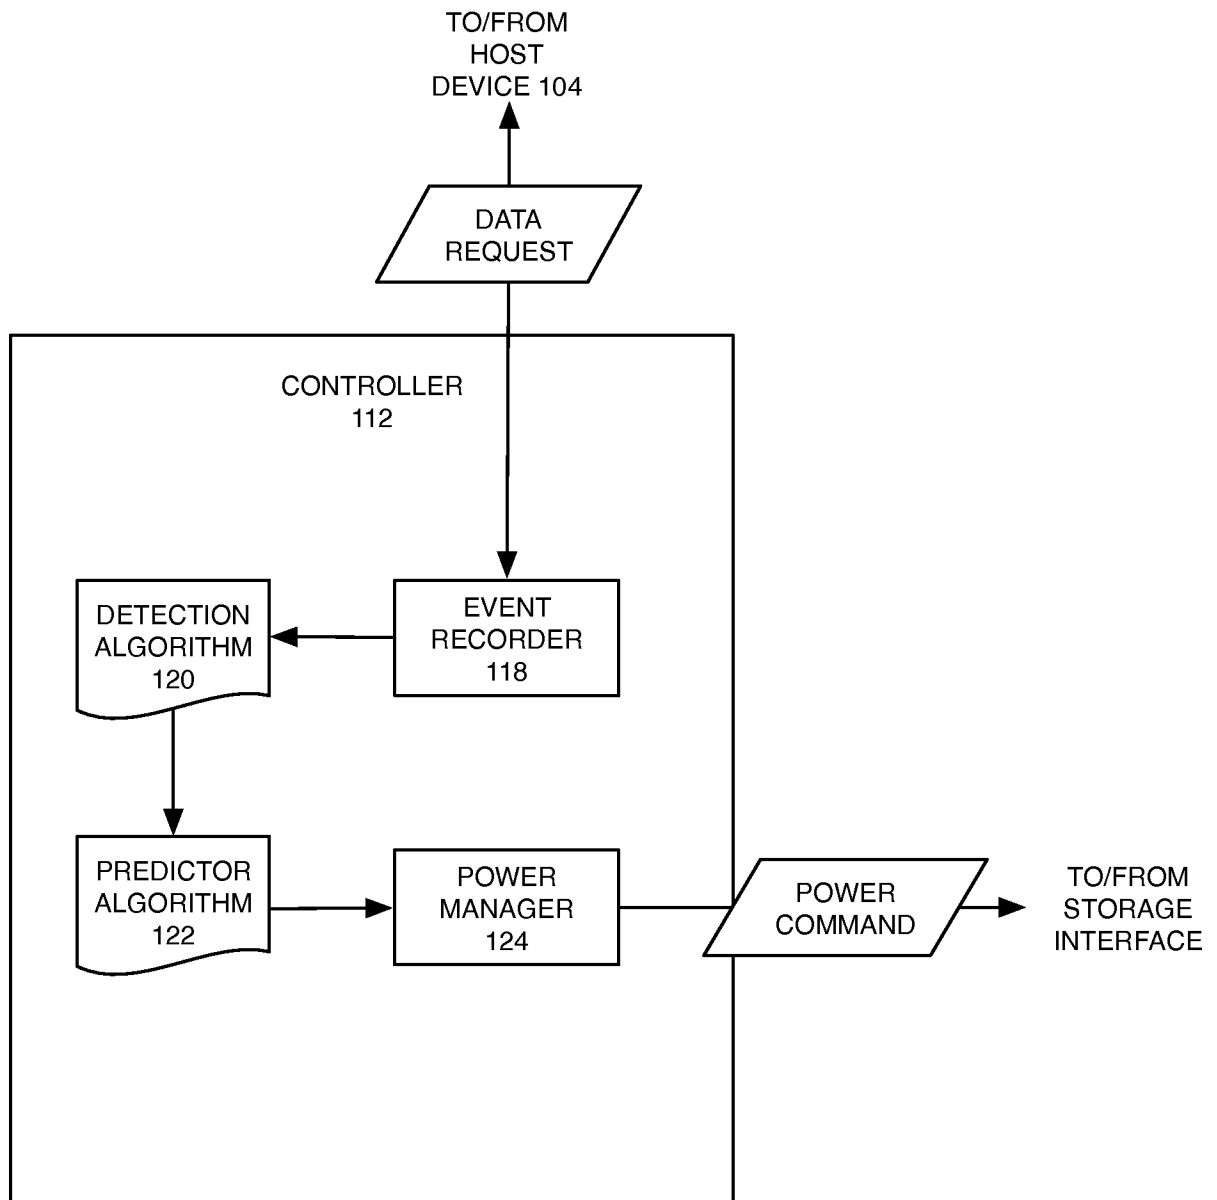


FIG. 1A

3/7




FIG. 2

4/7




FIG. 3



6/7




FIG. 5




FIG. 6

**A. CLASSIFICATION OF SUBJECT MATTER****G06F 19/00(2011.01)i**

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)

G06F 19/00; G06F 3/06; H04N 5/765; G06F 12/00; G06F 1/00; G06F 1/32

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) &amp; keywords: storage device, predict, future, access request, interval, pattern, HDD, standby

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

| Category* | Citation of document, with indication, where appropriate, of the relevant passages  | Relevant to claim No.   |
|-----------|---|---|
| Y<br>A    | US 2008-0104431 A1 (SHIMADA, KENTARO) 01 May 2008<br>See paragraphs [0039]-[0056], [0110]-[0123]; claims 1, 2, 10; and figures 1, 8-10. | 1-3, 8, 11-15, 18-20<br>, 22-25<br>4-7, 9, 10, 16, 17, 21<br>, 26, 27 |
| Y<br>A    | US 6892313 B1 (CODILIAN, RAFFI et al.) 10 May 2005<br>See abstract; column 2, line 57 - column 4, line 4; claim 1; and figures 1-3.     | 1-3, 8, 11-15, 18-20<br>, 22-25<br>4-7, 9, 10, 16, 17, 21<br>, 26, 27 |
| A         | US 2007-0150661 A1 (SUDA, YUKIHIRO) 28 June 2007<br>See paragraphs [0039]-[0076]; claims 1, 2; and figures 14-16.                       | 1-27  |
| A         | JP 2003-309797 A (SHARP CORP.) 31 October 2003<br>See abstract; paragraphs [0020]-[0024]; claim 1; and figure 2.                        | 1-27  |

 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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

02 September 2014 (02.09.2014)

Date of mailing of the international search report

**03 September 2014 (03.09.2014)**

Name and mailing address of the ISA/KR

International Application Division  
Korean Intellectual Property Office  
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,  
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

JANG, Ho Keun

Telephone No. +82-42-481-8187



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2014/032424**

| Patent document cited in search report | Publication date | Patent family member(s)   | Publication date                                     |
|--|------------------|---|--|
| US 2008-0104431 A1                     | 01/05/2008       | EP 1918809 A2<br>EP 1918809 A3<br>JP 2008-112292 A<br>US 7840820 B2 | 07/05/2008<br>07/07/2010<br>15/05/2008<br>23/11/2010 |
| US 6892313 B1                          | 10/05/2005       | None  |  |
| US 2007-0150661 A1                     | 28/06/2007       | JP 2007-179089 A  | 12/07/2007   |
| JP 2003-309797 A                       | 31/10/2003       | None  |  |