

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2012/148828 A2

(43) International Publication Date
1 November 2012 (01.11.2012)

(51) International Patent Classification:
G06F 12/02 (2006.01) *G11C 16/06* (2006.01)

(21) International Application Number:
PCT/US2012/034601

(22) International Filing Date:
22 April 2012 (22.04.2012)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/479,249 26 April 2011 (26.04.2011) US

(71) Applicant (for all designated States except US): **LSI CORPORATION** [US/US]; 1621 Barber Lane, Milpitas, California 95035 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **TOMLIN, Andrew, John** [GB/US]; 2824 Monte Cresta Way, San Jose, California 95132 (US).

(74) Agent: **SMITH, Walstein**; P.O. Box 1668, Georgetown, TX 78627-1668 (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, 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, IS, JP, KE, KG, KM, 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, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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, MD, 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, 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, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: VARIABLE OVER-PROVISIONING FOR NON-VOLATILE STORAGE




Fig. 3A

(57) Abstract: Dynamically varying Over-Provisioning (OP) enables improvements in lifetime, reliability, and/or performance of a Solid-State Disk (SSD) and/or a flash memory therein. A host coupled to the SSD writes newer data to the SSD. If the newer host data is less random than older host data, then entropy of host data on the SSD decreases. In response, an SSD controller of the SSD dynamically alters allocations of the flash memory, decreasing host allocation and increasing OP allocation. If the newer host data is more random, then the SSD controller dynamically increases the host allocation and decreases the OP allocation. The SSD controller dynamically allocates the OP allocation between host OP and system OP proportionally in accordance with a ratio of bandwidths of host and system data writes to the flash memory.

1 **VARIABLE OVER-PROVISIONING FOR NON-VOLATILE STORAGE**
2
3
45 **CROSS REFERENCE TO RELATED APPLICATIONS**
67 **[0001]** Priority benefit claims for this application are made in the accompanying
8 Application Data Sheet, Request, or Transmittal (as appropriate, if any). To the extent permitted
9 by the type of the instant application, this application incorporates by reference for all purposes
10 the following applications, all commonly owned with the instant application at the time the
11 invention was made:12 U.S. Provisional Application (Docket No. SF-11-04 and Serial No. 61/479,249), filed
13 04-26-2011, first named inventor Andrew John Tomlin, and entitled **Variable**
14 **Over-Provisioning for Non-Volatile Storage.**

15

16

17 **BACKGROUND**
1819 **[0002]** Field: Advancements in non-volatile storage technology are needed to provide
20 improvements in performance, efficiency, and utility of use.

21

22 **[0003]** Related Art: Unless expressly identified as being publicly or well known,
23 mention herein of techniques and concepts, including for context, definitions, or comparison
24 purposes, should not be construed as an admission that such techniques and concepts are
25 previously publicly known or otherwise part of the prior art. All references cited herein (if any),
26 including patents, patent applications, and publications, are hereby incorporated by reference in
27 their entireties, whether specifically incorporated or not, for all purposes.

28

29

30 **SYNOPSIS**
3132 **[0004]** The invention may be implemented in numerous ways, including as a process,
33 an article of manufacture, an apparatus, a system, a composition of matter, and a computer
34 readable medium such as a computer readable storage medium (e.g. media in an optical and/or
35 magnetic mass storage device such as a disk, or an integrated circuit having non-volatile storage
36 such as flash storage) or a computer network wherein program instructions are sent over optical

1 or electronic communication links. In this specification, these implementations, or any other
2 form that the invention may take, may be referred to as techniques. The Detailed Description
3 provides an exposition of one or more embodiments of the invention that enable improvements
4 in performance, efficiency, and utility of use in the field identified above. The Detailed
5 Description includes an Introduction to facilitate the more rapid understanding of the remainder
6 of the Detailed Description. The Introduction includes Example Embodiments of one or more of
7 systems, methods, articles of manufacture, and computer readable media in accordance with the
8 concepts described herein. As is discussed in more detail in the Conclusions, the invention
9 encompasses all possible modifications and variations within the scope of the issued claims.

10

1 Brief Description of Drawings

[0005] Fig. 1A illustrates selected details of an embodiment of a Solid-State Disk (SSD) including an SSD controller using variable Over-Provisioning (OP) for managing non-volatile storage, such as implemented via Non-Volatile Memory (NVM) elements (e.g. flash memories).

[0006] Fig. 1B illustrates selected details of various embodiments of systems including one or more instances of the SSD of Fig. 1A.

10
11 [0007] Fig. 2 illustrates selected details of host and system allocations of flash memory,
12 including for OP use, according to various embodiments of a system using variable OP for
13 managing NVMs.

[0008] Fig. 3A and Fig. 3B illustrate selected details of various embodiments of dynamically varying host and system OPs as relating to dynamically varying data entropy.

18 [0009] Fig. 4 illustrates a flow diagram of an embodiment of (re)allocation of OP
19 resources in a context of variable OP for managing NVMs.

21 [0010] Fig. 5 illustrates a flow diagram of an embodiment of a life-cycle of a managed
22 unit of flash memory.

1 List of Reference Symbols in Drawings

2

3 [0011]

4

Ref. Symbol	Element Name
100	SSD Controller
101	SSD
102	Host
103	(optional) Switch / Fabric / Intermediate Controller
104	Intermediate Interfaces
105	OS
106	FirmWare (FW)
107	Driver
107D	dotted-arrow (Host Software \leftrightarrow I/O Device Communication)
109	Application
109D	dotted-arrow (Application \leftrightarrow I/O Device Communication via driver)
109V	dotted-arrow (Application \leftrightarrow I/O Device Communication via VF)
110	External Interfaces
111	Host Interfaces
112C	(optional) Card Memory
113	Tag Tracking
114	Multi-Device Management Software
115	Host Software
116	I/O Card
117	I/O & Storage Devices/Resources
118	Servers
119	LAN/WAN
121	Data Processing
123	Engines
131	Buffer
133	DMA
135	ECC-X
137	Memory
141	Map
143	Table

Ref. Symbol	Element Name
151	Recycler
161	ECC
171	CPU
172	CPU Core
173	Command Management
175	Buffer Management
177	Translation Management
179	Coherency Management
180	Memory Interface
181	Device Management
182	Identity Management
190	Device Interfaces
191	Device Interface Logic
192	Flash Device
193	Scheduling
194	Flash Die
199	NVM
200	Flash Memory
201	System Allocation
202A	Host Allocation
202C	Host Allocation
203A	OP Allocation
203C	OP Allocation
204A	System OP Allocation
204B	System OP Allocation
204C	System OP Allocation
205A	Host OP Allocation
205B	Host OP Allocation
208	System OP Allocation Delta
209	System OP Allocation Delta
210A	Initial Allocation
210B	Same Allocation
210C	Increased Allocation
301	System Allocation

Ref. Symbol	Element Name
301F	System Allocation (fixed)
302G1	Host Allocation
302G2	Host Allocation
302H1	Host Allocation
302H2	Host Allocation
302V	Host Allocation
303G1	(combined) OP Allocation
303G2	(combined) OP Allocation
303H1	(combined) OP Allocation
303H2	(combined) OP Allocation
303V1	System OP Allocation (var)
303V2	System OP Allocation (var)
304F	System OP Allocation (fixed)
304G1	System OP Allocation
304G2	System OP Allocation
304H1	System OP Allocation
304H2	System OP Allocation
305G1	Host OP Allocation
305G2	Host OP Allocation
305H1	Host OP Allocation
305H2	Host OP Allocation
310	Physical Size
320	Data Entropy
330G	G%
330H	H%
400	Flow Diagram
401	Determine If (Re)Allocation Conditions Exist
402	(Re)Allocate?
403	Perform (Re)Allocation
500	Flow Diagram
501	Free
502	Free Queue
503	Host Unit
504	System Unit

Ref. Symbol	Element Name
512	Queue Allocation
513	Host Allocation
514	System Allocation
515	Recycle Host Unit
516	Recycle System Unit

DETAILED DESCRIPTION

2 [0012] A detailed description of one or more embodiments of the invention is provided
3 below along with accompanying figures illustrating selected details of the invention. The
4 invention is described in connection with the embodiments. The embodiments herein are
5 understood to be merely exemplary, the invention is expressly not limited to or by any or all of
6 the embodiments herein, and the invention encompasses numerous alternatives, modifications,
7 and equivalents. To avoid monotony in the exposition, a variety of word labels (including but
8 not limited to: first, last, certain, various, further, other, particular, select, some, and notable)
9 may be applied to separate sets of embodiments; as used herein such labels are expressly not
10 meant to convey quality, or any form of preference or prejudice, but merely to conveniently
11 distinguish among the separate sets. The order of some operations of disclosed processes is
12 alterable within the scope of the invention. Wherever multiple embodiments serve to describe
13 variations in process, method, and/or program instruction features, other embodiments are
14 contemplated that in accordance with a predetermined or a dynamically determined criterion
15 perform static and/or dynamic selection of one of a plurality of modes of operation
16 corresponding respectively to a plurality of the multiple embodiments. Numerous specific
17 details are set forth in the following description to provide a thorough understanding of the
18 invention. The details are provided for the purpose of example and the invention may be
19 practiced according to the claims without some or all of the details. For the purpose of clarity,
20 technical material that is known in the technical fields related to the invention has not been
21 described in detail so that the invention is not unnecessarily obscured.

24 INTRODUCTION

26 [0013] This introduction is included only to facilitate the more rapid understanding of
27 the Detailed Description; the invention is not limited to the concepts presented in the
28 introduction (including explicit examples, if any), as the paragraphs of any introduction are
29 necessarily an abridged view of the entire subject and are not meant to be an exhaustive or
30 restrictive description. For example, the introduction that follows provides overview
31 information limited by space and organization to only certain embodiments. There are many
32 other embodiments, including those to which claims will ultimately be drawn, discussed
33 throughout the balance of the specification.

1 Acronyms

2

3 [0014] At least some of the various shorthand abbreviations (e.g. acronyms) defined
 4 here refer to certain elements used herein.

<u>Acronym</u>	<u>Description</u>
AHCI	Advanced Host Controller Interface
API	Application Program Interface
BCH	Bose Chaudhuri Hocquenghem
ATA	Advanced Technology Attachment (AT Attachment)
CD	Compact Disk
CF	Compact Flash
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
DAS	Direct Attached Storage
DDR	Double-Data-Rate
DES	Data Encryption Standard
DMA	Direct Memory Access
DNA	Direct NAND Access
DRAM	Dynamic Random Access Memory
DVD	Digital Versatile/Video Disk
DVR	Digital Video Recorder
ECC	Error-Correcting Code
eMMC	Embedded MultiMediaCard
eSATA	external Serial Advanced Technology Attachment
GPS	Global Positioning System
I/O	Input / Output
IC	Integrated Circuit
IDE	Integrated Drive Electronics
LAN	Local Area Network
LB	Logical Block
LBA	Logical Block Address
LDPC	Low-Density Parity-Check
LPN	Logical Page Number
MLC	Multi-Level Cell
MMC	MultiMediaCard

<u>Acronym</u>	<u>Description</u>
NAS	Network Attached Storage
NCQ	Native Command Queuing
NVM	Non-Volatile Memory
ONA	Optimized NAND Access
ONFI	Open NAND Flash Interface
OP	Over-Provisioning
OS	Operating System
PC	Personal Computer
PCIe	Peripheral Component Interconnect express (PCI express)
PDA	Personal Digital Assistant
POS	Point Of Sale
RAID	Redundant Array of Inexpensive/Independent Disks
ReRAM	Resistive Random Access Memory
RS	Reed-Solomon
SAN	Storage Attached Network
SAS	Serial Attached Small Computer System Interface (Serial SCSI)
SATA	Serial Advanced Technology Attachment (Serial ATA)
SCSI	Small Computer System Interface
SD	Secure Digital
SDR	Single-Data-Rate
SLC	Single-Level Cell
SMART	Self-Monitoring Analysis and Reporting Technology
SSD	Solid-State Disk/Drive
UFS	Unified Flash Storage
USB	Universal Serial Bus
VF	Virtual Function
WAN	Wide Area Network

1

2

3

1 [0015] Conceptually, an SSD controller providing variable OP for NVM, such as used
2 for non-volatile storage of SSD data, enables longer lifetimes, enhanced reliability, and/or
3 improved performance, in various circumstances. Conceptually, the NVM is dynamically
4 allocated between storing host data, storing system data, and OP use, and the OP use allocation
5 is dynamically allocated between host data OP and system data OP. The allocations are
6 dynamically variable based on various operating characteristics and/or contexts of the SSD, the
7 SSD controller, and/or the NVM.

8

9 [0016] A host coupled to the SSD writes newer data to the SSD. If the newer host data
10 is less random than older host data (and thus more effectively compressible), then entropy of the
11 host data is decreasing. In response, an SSD controller of the SSD dynamically alters allocations
12 of the NVM (e.g. flash memory), decreasing host allocation while increasing OP allocation, and
13 dynamically allocates the OP allocation between system OP and host OP proportionally in
14 accordance with a ratio of bandwidths of system and host data writes to the NVM. If the newer
15 host data is more random (and thus less effectively compressible), then the host allocation is
16 increased, the OP allocation is decreased, and allocated between the system and the host OPs in
17 accordance with the ratio.

18

19 [0017] For instance, as “free space” of the NVM increases/decreases (e.g. due to
20 reduced/increased host or system data storage space), the OP use allocation is
21 increased/decreased (optionally after a delay to accommodate garbage collection, recycling,
22 and/or erasure). For another instance, the OP use allocation between system data OP and host
23 data OP is proportional to a dynamically varying value, the dynamically varying value being a
24 bandwidth of system data writes to the NVM divided by a bandwidth of host data writes to the
25 NVM. The host data storage space varies at least according to varying entropy of the host data
26 (e.g. compressibility and/or responsiveness to size reduction via de-duplication) and host
27 commands that explicitly de-allocate previously stored host data.

28

29 [0018] For a first example, if the NVM available for OP decreases (increases), by an
30 amount, then the system data OP and the host data OP are collectively decreased (increased) by
31 the amount while maintaining a same ratio between the system data OP and the host data OP.
32 For a second example, if a ratio of a system data rate to a host data rate changes, then the system
33 data OP and the host data OP are adjusted according to the change in the ratio. An instance of a
34 system data rate is a bandwidth of system data writes to NVM, and an instance of a host data
35 rate is a bandwidth of host data writes (e.g. writes that are not system data writes) to NVM, such
36 that the system data writes correspond to all NVM writes except NVM writes that are writing

1 host data. The writing of host data optionally and/or selectively includes writing headers usable
2 to identify the host data and/or ECC information corresponding to the host data. As a third
3 example, system data OP and host data OP are allocated in accordance with respective data rates
4 (e.g. linearly in accordance with a ratio of the data rates), such as when a unit of NVM is
5 allocated, or when a garbage collection (e.g. to reclaim partially used NVM units) is performed.
6

7 **[0019]** For a fourth example, the first and the second (or third) examples are combined,
8 such that a change in NVM available for OP and a change in a system to host data rate ratio
9 result in corresponding changes to system data OP and host data OP allocations. Additional
10 examples include any of the aforementioned examples where the system data OP and/or the host
11 data OP are reallocated dynamically based on an event, such as an event associated with
12 allocation of a unit of the NVM, garbage collection of a portion of the NVM, or any other event
13 where a unit of the NVM is available for reallocation from one type of use (e.g. host data) to
14 another type of use (e.g. system data).

15

16 **[0020]** In some embodiments, allocation of OP resources between system and host
17 usages is subject to respective minimum/maximum values and/or granularities. For example,
18 system and/or host OP allocation is a minimum of a predetermined number of units, independent
19 of host allocation and/or ratio of system data rate to host data rate. For another example, system
20 and/or host OP allocation is granular in accordance with a predetermined number of units.

21

22 **[0021]** In various embodiments, NVM (e.g. flash memory) is managed in portions,
23 referred to as (managed) units of (flash memory), and also referred to herein as 'managed units'
24 or 'units'. Examples of (managed) units of (flash memory) are one or more contiguous and/or
25 non-contiguous portions of the flash memory, such as one or more contiguous/non-contiguous
26 pages/blocks of the flash memory, one or more R-blocks (described elsewhere herein) of the
27 flash memory, or any sub-portion of the flash memory suitable for management operations (such
28 as allocation). In some embodiments, allocation of flash memory is granular in accordance with
29 integer quanta of flash memory management units.

30

31 **[0022]** An example of an R-block is a logical slice or section across all die of a flash
32 memory. For example, in a flash memory having R flash die, each flash die having N blocks,
33 each R-block is the i^{th} block from each of the flash die taken together, for a total of N R-blocks.
34 For another example, in a flash memory having R flash die, each with N blocks, each R-block is
35 the i^{th} and $(i+1)^{\text{th}}$ block from each of the flash die, for a total of $N/2$ R-blocks. For yet another

1 example, in a flash memory having a plurality of dual plane devices, each R-block is the i^{th} even
2 block and the i^{th} odd block from each of the dual plane devices.

3

4 [0023] In some situations, write amplification results when a host storage write of a
5 particular size results in a plurality of writes (each having a size of, e.g., a multiple of the
6 particular size) to flash memory of an SSD. The plurality of writes arises from, for example,
7 erasing a portion of flash memory before writing (e.g. programming) the portion, wear leveling,
8 garbage collection, and flash memory management operations that result in system data writes.
9 An example calculation of write amplification is an amount of data written to flash memory on
10 behalf of a particular collection of host writes (including, e.g., system writes to complete writing
11 of host data associated with the host writes), divided by an amount of data written by the
12 particular collection of host writes.

13

14 [0024] In some embodiments and/or usage scenarios, write amplification is relatively
15 low, such as a relatively long sequence of host writes to sequential addresses when garbage
16 collection and/or other operations have made a relatively large portion of flash memory erased
17 and readily available for writing (e.g. programming). In some embodiments and/or usage
18 scenarios, write amplification is relatively high, such as a relatively long sequence of host writes
19 to random addresses when a relatively small portion of flash memory is erased and readily
20 available for writing.

21

22 [0025] Via a controller for an SSD, flash memory of the SSD is allocated between
23 storage of host data and system data. The host data is associated with a host coupled to the SSD
24 and the system data is associated with the controller. The controller dynamically controls the
25 allocation according to overall free space of the flash memory. In some embodiments and/or
26 usage scenarios, increasing OP of the system data via increasing the allocation to the system data
27 (and decreasing the allocation to the host data) enables reduced write amplification and
28 improved SSD and/or flash memory performance, reliability, and/or lifetime. The overall free
29 space of the flash memory is the free space of the host data allocation and the free space of the
30 system data allocation.

31

32 [0026] The dynamic allocation control is used in various embodiments, for instance
33 where the host data includes logical units of various sizes, such as when the host data is
34 compressed, deduplicated, or transformed in some manner resulting in logical units of various
35 sizes. The dynamic allocation control is used in various embodiments, for instance where the

1 host data is relatively highly reducible, such as when the host data is relatively highly
2 compressible or relatively highly compactable via deduplication.

3

4 **[0027]** The host data includes data received from (or provided to) one or more hosts
5 (e.g. computing hosts), or data derived from the data from (or to) the hosts (e.g. via compression,
6 deduplication, encryption, or any reversible transformation). Examples of host data include user
7 data and/or files, application data and/or files, driver data and/or files, OS data, data structures,
8 and/or files, and other information provided by the host via an interface between the SSD and
9 the host. The host data as written to NVM optionally and/or selectively includes metadata added
10 by the SSD controller and written with the host data, such as headers used to identify the host
11 data and/or ECC information corresponding to the host data. The system data includes data
12 relating to management or operation of the controller, the SSD, and/or the flash memory, or any
13 portions thereof. Examples of system data include LBA to flash memory block mapping
14 information and checkpoint information to enable restoration of correct state in an event of a
15 malfunction (e.g. a crash or a power failure). Other examples of system data include
16 information not specific to host data but useful in managing storage thereof via an SSD, SSD
17 controller, and/or NVM (e.g. flash memory), or managing operation of an SSD, SSD controller,
18 and/or NVM.

19

20 **[0028]** In various situations, a ratio of host data space used to system data space used is
21 a ratio of 100:1 to 500:1. In various situations, such as situations with a relatively high number
22 of random writes, a ratio of system data write bandwidth to host data write bandwidth is a ratio
23 of 2:1 to 5:1.

24

25 **[0029]** Elsewhere herein, bandwidth of host data writes to flash memory is sometimes
26 referred to as bandwidth of host data writes or alternatively as host write bandwidth and
27 bandwidth of system data writes to flash memory is sometimes referred to as bandwidth of
28 system data writes or alternatively as system write bandwidth.

29

30 **[0030]** In some embodiments, accessing compressed data of varying-sized quanta in
31 NVM provides improved storage efficiency in some usage scenarios. For example, an SSD
32 controller receives data (that is not compressed) from a computing host (e.g., relating to a disk
33 write command), compresses the data, and stores the compressed data into flash memory. In
34 response to a subsequent request from the computing host (e.g., relating to a disk read
35 command), the SSD controller reads the compressed data from the flash memory, decompresses
36 the compressed data, and provides the decompressed data to the computing host. The

1 compressed data is stored in the flash memory according to varying-sized quanta, the quanta size
2 varying due to, e.g., compression algorithm, operating mode, and compression effectiveness on
3 various data. The SSD controller decompresses the data in part by consulting an included map
4 table to determine where header(s) are stored in the flash memory. The SSD controller parses
5 the header(s) obtained from the flash memory to determine where appropriate (compressed) data
6 is stored in the flash memory. The SSD controller decompresses the appropriate data from the
7 flash memory to produce the decompressed data to provide to the computing host.

8

9 **[0031]** In various embodiments, an SSD controller includes a host interface for
10 interfacing with a computing host, an interface for interfacing with NVM such as flash memory,
11 and circuitry for controlling the interfaces and performing (and/or controlling various aspects of
12 the performing) compressing and decompressing, as well as lower-level error correction, higher-
13 level error correction, and dynamic higher-level redundancy mode management with
14 independent silicon elements.

15

16 **[0032]** According to various embodiments, some host interfaces are compatible with
17 one or more of a USB interface standard, a CF interface standard, an MMC interface standard,
18 an eMMC interface standard, a Thunderbolt interface standard, a UFS interface standard, an SD
19 interface standard, a Memory Stick interface standard, an xD-picture card interface standard, an
20 IDE interface standard, a SATA interface standard, a SCSI interface standard, a SAS interface
21 standard, and a PCIe interface standard. According to various embodiments, the computing host
22 is all or any portions of a computer, a workstation computer, a server computer, a storage server,
23 a SAN, a NAS device, a DAS device, a storage appliance, a PC, a laptop computer, a notebook
24 computer, a netbook computer, a tablet device or computer, an ultrabook computer, an electronic
25 reading device (such as an e-reader), a PDA, a navigation system, a (handheld) GPS device, an
26 automotive control system, an automotive media control system or computer, a printer, copier or
27 fax machine or all-in-one device, a POS device, a cash-register, a media player, a television, a
28 media recorder, a DVR, a digital camera, a cellular handset, a cordless telephone handset, and an
29 electronic game. In some embodiments, an interfacing host (such as an SAS/SATA bridge)
30 operates as a computing host and/or as a bridge to a computing host.

31

32 **[0033]** In various embodiments, the SSD controller includes one or more processors.
33 The processors execute firmware to control and/or perform operation of the SSD controller. The
34 SSD controller communicates with the computing host to send and receive commands and/or
35 status as well as data. The computing host executes one or more of an operating system, a
36 driver, and an application. Communication by the computing host with the SSD controller is

1 optionally and/or selectively via the driver and/or via the application. In a first example, all
2 communication to the SSD controller is via the driver, and the application provides higher-level
3 commands to the driver that the driver translates into specific commands for the SSD controller.
4 In a second example, the driver implements a bypass mode and the application is enabled to send
5 specific commands to the SSD controller via the driver. In a third example, a PCIe SSD
6 controller supports one or more Virtual Functions (VFs), enabling an application, once
7 configured, to communicate directly with the SSD controller, bypassing the driver.
8

9 **[0034]** According to various embodiments, some SSDs are compatible with form-
10 factors, electrical interfaces, and/or protocols used by magnetic and/or optical non-volatile
11 storage, such as HDDs, CD drives, and DVD drives. In various embodiments, SSDs use various
12 combinations of zero or more parity codes, zero or more RS codes, zero or more BCH codes,
13 zero or more Viterbi or other trellis codes, and zero or more LDPC codes.

14

15

16 EXAMPLE EMBODIMENTS

17

18 **[0035]** In concluding the introduction to the detailed description, what follows is a
19 collection of example embodiments, including at least some explicitly enumerated as “ECs”
20 (Example Combinations), providing additional description of a variety of embodiment types in
21 accordance with the concepts described herein; these examples are not meant to be mutually
22 exclusive, exhaustive, or restrictive; and the invention is not limited to these example
23 embodiments but rather encompasses all possible modifications and variations within the scope
24 of the issued claims and their equivalents.

25

26 **[0036]** EC1) A system comprising:

27 a means for operating all or any portions of one or more flash memories as respective
28 allocations dedicated respectively to host data, system data, system over-
29 provisioning (OP), and host OP; and
30 a means for dynamically determining any one or more of the respective allocations in
31 response to one or more events.

32

33 **[0037]** EC2) The system of EC1, wherein the events comprise a change in amount of
34 usage of any one or more of the respective allocations.

35

1 [0038] EC3) The system of EC1, wherein the events comprise a determination of
2 altered effectiveness of any one or more of compression, deduplication, and transformation of
3 information stored in the flash memories.

4

5 [0039] EC4) The system of EC3, wherein the altered effectiveness is measured at least
6 in part by a particular amount of space used in a particular one of the respective allocations.

7

8 [0040] EC5) The system of EC3, wherein the altered effectiveness is measured at least
9 in part by a particular amount of free space available in a particular one of the respective
10 allocations.

11

12 [0041] EC6) The system of EC1, wherein the events comprise a determination of
13 increased effectiveness of any one or more of compression and deduplication of information
14 stored in the host data allocation.

15

16 [0042] EC7) The system of EC6, wherein the increased effectiveness is measured at
17 least in part by a reduction in amount of space used in the host data allocation.

18

19 [0043] EC8) The system of EC6, wherein the increased effectiveness is measured at
20 least in part by an increase in amount of free space available in the host data allocation.

21

22 [0044] EC9) The system of EC6, wherein the means for dynamically determining
23 dynamically reduces any one or more of the host data allocation and the host OP allocation in
24 response to the determination of increased effectiveness.

25

26 [0045] EC10) The system of EC6, wherein the means for dynamically determining
27 dynamically increases any one or more of the system data allocation and the system OP
28 allocation in response to the determination of increased effectiveness.

29

30 [0046] EC11) The system of EC1, wherein the means for operating and the means for
31 dynamically determining are comprised in a controller of a Solid-State Disk (SSD) and the flash
32 memories are comprised in the SSD.

33

1 [0047] EC12) A method comprising:
2 operating one or more flash memories of a Solid-State Disk (SSD) as respective host
3 data, system data, and combined Over-Provisioning (OP) allocations;
4 operating the combined OP allocation as respective system and host OP allocations;
5 dynamically determining a new value of the system OP allocation as a linear function of
6 a ratio of a system data rate to a host data rate;
7 dynamically altering the system OP allocation in response to changes in the new value;
8 and
9 wherein the dynamically determining is in response to an event.

10

11 [0048] EC13) The method of EC12, further comprising.
12 receiving data from a host coupled to the SSD;
13 determining information from the data received from the host;
14 storing the information into the flash memories in accordance with the host data
15 allocation.

16

17 [0049] EC14) The method of EC13, further comprising computing the information via
18 compressing the data received from the host.

19

20 [0050] EC15) The method of EC13, further comprising computing the information via
21 deduplicating the data received from the host.

22

23 [0051] EC16) The method of EC12, further comprising storing host address to flash
24 memory address correspondence information into the flash memories in accordance with the
25 system data allocation.

26

27 [0052] EC17) The method of EC12, wherein the event is based at least in part on
28 selecting a unit of the flash memories for reclamation.

29

30 [0053] EC18) The method of EC17, wherein the reclamation comprises one or more of.

1 garbage collection,
2 selection for garbage collection,
3 recycling,
4 selection for recycling,
5 erasure, and
6 selection for erasure.

7

8 **[0054]** EC19) A system comprising:
9 a means for operating one or more flash memories as respective host data, system data,
10 and combined over-provisioning (OP) allocations;
11 a means for operating the combined OP allocation as respective system and host OP
12 allocations; and
13 a means for dynamically determining a new value of the system OP allocation in
14 response to one or more events.

15

16 **[0055]** EC20) The system of EC19, further comprising a means for dynamically
17 altering the system OP allocation in response to changes in the new value.

18

19 **[0056]** EC21) The system of EC19, wherein the means for dynamically determining the
20 new value determines the new value in accordance with a function of a system data rate and a
21 host data rate.

22

23 **[0057]** EC22) The system of EC21, wherein the means for dynamically determining the
24 new value determines the new value in accordance with a ratio of the system data rate to the host
25 data rate.

26

27 **[0058]** EC23) The system of EC22, wherein the means for dynamically determining the
28 new value determines the new value as being proportional to the system data rate divided by the
29 host data rate.

30

31 **[0059]** EC24) The system of EC21, wherein the means are via a controller of a Solid-
32 State Disk (SSD) and the flash memories are comprised in the SSD.

33

1 [0060] EC25) The system of EC24, further comprising a means for interfacing to a host
2 and wherein the means for interfacing to the host is comprised in the SSD and the host data
3 allocation is used at least in part to store information representing at least a portion of data
4 communicated via the means for interfacing to the host.

5

6 [0061] EC26) The system of EC25, further comprising a means for determining at least
7 a portion of the information via compression of at least a portion of the data communicated.

8

9 [0062] EC27) The system of EC25, further comprising a means for determining at least
10 a portion of the information via deduplication of at least a portion of the data communicated.

11

12 [0063] EC28) The system of EC25, further comprising a means for determining at least
13 a portion of the information via a reversible transformation that decreases entropy of at least a
14 portion of the data communicated.

15

16 [0064] EC29) The system of EC24, wherein the system data allocation is used at least
17 in part by the controller to store at least a portion of map data.

18

19 [0065] EC30) The system of EC24, wherein the system data allocation is used at least
20 in part by the controller to store at least a portion of checkpoint data.

21

22 [0066] EC31) The system of EC21, wherein the system data rate is based at least in part
23 on one or more of
24 a rate of system data provided from a Solid-State-Disk (SSD) controller to store in the
25 flash memories,
26 a rate of system data written to the flash memories, and
27 a bandwidth of system data writes to the flash memories.

28

29 [0067] EC32) The system of EC21, wherein the host data rate is based at least in part
30 on one or more of
31 a rate of data provided from a host to store in the flash memories,
32 a rate of host data written to the flash memories, and
33 a bandwidth of host data writes to the flash memories.

34

35 [0068] EC33) The system of EC19, wherein the events comprise initiating, completing,
36 requesting, and/or selecting for one or more of

1 allocation of a unit of the flash memories,
2 garbage collection of one or more portions of the flash memories,
3 recycling of one or more portions of the flash memories, and
4 erasure of one or more portions of the flash memories.

5

6 [0069] EC34) The system of EC19, further comprising a means for communicating at
7 least a portion of host data represented by information stored in at least a part of the host data
8 allocation at least in part in response to requests from a computing host.

9

10 [0070] EC35) The system of EC34, further comprising a means for interfacing the
11 requests with the computing host.

12

13 [0071] EC36) The system of EC35, wherein the means for interfacing the requests with
14 the computing host is compatible with a storage interface standard.

15

16 [0072] EC37) The system of EC36, wherein the storage interface standard comprises
17 one or more of

18 a Universal Serial Bus (USB) interface standard,
19 a Compact Flash (CF) interface standard,
20 a MultiMediaCard (MMC) interface standard,
21 an embedded MMC (eMMC) interface standard,
22 a Thunderbolt interface standard,
23 a UFS interface standard,
24 a Secure Digital (SD) interface standard,
25 a Memory Stick interface standard,
26 an xD-picture card interface standard,
27 an Integrated Drive Electronics (IDE) interface standard,
28 a Serial Advanced Technology Attachment (SATA) interface standard,
29 an external SATA (eSATA) interface standard,
30 a Small Computer System Interface (SCSI) interface standard,
31 a Serial Attached Small Computer System Interface (SAS) interface standard,
32 a Fibre Channel interface standard,
33 an Ethernet interface standard, and
34 a Peripheral Component Interconnect express (PCIe) interface standard.

35

1 [0073] EC38) The system of EC34, further comprising all or any portions of the
2 computing host.
3

1 [0074] EC39) The system of EC38, wherein the computing host comprises one or more
2 of

3 a computer,
4 a workstation computer,
5 a server computer,
6 a storage server,
7 a Storage Attached Network (SAN),
8 a Network Attached Storage (NAS) device,
9 a Direct Attached Storage (DAS) device,
10 a storage appliance,
11 a Personal Computer (PC),
12 a laptop computer,
13 a notebook computer,
14 a netbook computer,
15 a tablet device or computer,
16 an ultrabook computer,
17 an electronic reading device (an e-reader),
18 a Personal Digital Assistant (PDA),
19 a navigation system,
20 a (handheld) Global Positioning System (GPS) device,
21 an automotive control system,
22 an automotive media control system or computer,
23 a printer, copier or fax machine or all-in-one device,
24 a Point Of Sale POS device,
25 a cash-register,
26 a media player,
27 a television,
28 a media recorder,
29 a Digital Video Recorder (DVR),
30 a digital camera,
31 a cellular handset,
32 a cordless telephone handset, and
33 an electronic game.

34

35 [0075] EC40) The system of EC19, further comprising a means for interfacing with the
36 flash memories.

1

2 [0076] EC41) The system of EC40, wherein the means for interfacing with the flash
3 memories comprises a flash memory interface.

4

5 [0077] EC42) The system of EC41, wherein the flash memory interface is compatible
6 with one or more of
7 an Open NAND Flash Interface (ONFI),
8 a Toggle-mode interface,
9 a Double-Data-Rate (DDR) synchronous interface,
10 a DDR2 synchronous interface;
11 a synchronous interface, and
12 an asynchronous interface.

13

14 [0078] EC43) The system of EC19, further comprising at least one of the flash
15 memories.

16

17 [0079] EC44) The system of EC43, wherein the at least one flash memory comprises
18 one or more of
19 NAND flash technology storage cells, and
20 NOR flash technology storage cells.

21

22 [0080] EC45) The system of EC43, wherein the at least one flash memory comprises
23 one or more of
24 Single-Level Cell (SLC) flash technology storage cells, and
25 Multi-Level Cell (MLC) flash technology storage cells.

26

27 [0081] EC46) The system of EC43, wherein the at least one flash memory comprises
28 one or more of
29 polysilicon technology-based charge storage cells, and
30 silicon nitride technology-based charge storage cells.

31

32 [0082] EC47) The system of EC43, wherein the at least one flash memory comprises
33 one or more of
34 two-dimensional technology-based flash memory technology, and
35 three-dimensional technology-based flash memory technology.

36

1 [0083] EC48) The system of EC19, further comprising:
2 a means for interfacing requests from a computing host, the requests relating to
3 information stored in the flash memories; and
4 a means for interfacing to the flash memories.
5

6 [0084] EC49) The system of EC48, wherein the means are collectively implemented in
7 a single Integrated Circuit (IC).
8

9 [0085] EC50) The system of EC48, wherein the means are comprised in a Solid-State
10 Disk (SSD).
11

1 SYSTEM

2

3 [0086] Fig. 1A illustrates selected details of an embodiment of SSD **101** including an
4 SSD controller using variable OP for managing non-volatile storage, such as implemented via
5 NVM elements (e.g. flash memories). The SSD controller is for managing non-volatile storage,
6 such as implemented via NVM elements (e.g., flash memories). SSD Controller **100** is
7 communicatively coupled via one or more External Interfaces **110** to a host (not illustrated).
8 According to various embodiments, External Interfaces **110** are one or more of: a SATA
9 interface; a SAS interface; a PCIe interface; a Fibre Channel interface; an Ethernet Interface
10 (such as 10 Gigabit Ethernet); a non-standard version of any of the preceding interfaces; a
11 custom interface; or any other type of interface used to interconnect storage and/or
12 communications and/or computing devices. For example, in some embodiments, SSD
13 Controller **100** includes a SATA interface and a PCIe interface.

14

15 [0087] SSD Controller **100** is further communicatively coupled via one or more Device
16 Interfaces **190** to NVM **199** including one or more storage devices, such as one or more of Flash
17 Device **192**. According to various embodiments, Device Interfaces **190** are one or more of: an
18 asynchronous interface; a synchronous interface; a single-data-rate (SDR) interface; a double-
19 data-rate (DDR) interface; a DRAM-compatible DDR or DDR2 synchronous interface; an ONFI
20 compatible interface, such as an ONFI 2.2 or ONFI 3.0 compatible interface; a Toggle-mode
21 compatible flash interface; a non-standard version of any of the preceding interfaces; a custom
22 interface; or any other type of interface used to connect to storage devices.

23

24 [0088] Each Flash Device **192** has, in some embodiments, one or more individual Flash
25 Die **194**. According to type of a particular one of Flash Device **192**, a plurality of Flash Die **194**
26 in the particular Flash Device **192** is optionally and/or selectively accessible in parallel. Flash
27 Device **192** is merely representative of one type of storage device enabled to communicatively
28 couple to SSD Controller **100**. In various embodiments, any type of storage device is usable,
29 such as an SLC NAND flash memory, MLC NAND flash memory, NOR flash memory, flash
30 memory using polysilicon or silicon nitride technology-based charge storage cells, two- or three-
31 dimensional technology-based flash memory, read-only memory, static random access memory,
32 dynamic random access memory, ferromagnetic memory, phase-change memory, racetrack
33 memory, ReRAM, or any other type of memory device or storage medium.

34

35 [0089] According to various embodiments, Device Interfaces **190** are organized as:
36 one or more busses with one or more of Flash Device **192** per bus; one or more groups of busses

1 with one or more of Flash Device **192** per bus, where busses in a group are generally accessed in
2 parallel; or any other organization of one or more of Flash Device **192** onto Device Interfaces
3 **190**.

4

5 **[0090]** Continuing in Fig. 1A, SSD Controller **100** has one or more modules, such as
6 Host Interfaces **111**, Data Processing **121**, Buffer **131**, Map **141**, Recycler **151**, ECC **161**, Device
7 Interface Logic **191**, and CPU **171**. The specific modules and interconnections illustrated in Fig.
8 1A are merely representative of one embodiment, and many arrangements and interconnections
9 of some or all of the modules, as well as additional modules not illustrated, are conceived. In a
10 first example, in some embodiments, there are two or more Host Interfaces **111** to provide dual-
11 porting. In a second example, in some embodiments, Data Processing **121** and/or ECC **161** are
12 combined with Buffer **131**. In a third example, in some embodiments, Host Interfaces **111** is
13 directly coupled to Buffer **131**, and Data Processing **121** optionally and/or selectively operates
14 on data stored in Buffer **131**. In a fourth example, in some embodiments, Device Interface Logic
15 **191** is directly coupled to Buffer **131**, and ECC **161** optionally and/or selectively operates on
16 data stored in Buffer **131**.

17

18 **[0091]** Host Interfaces **111** sends and receives commands and/or data via External
19 Interfaces **110**, and, in some embodiments, tracks progress of individual commands via Tag
20 Tracking **113**. For example, the commands include a read command specifying an address (such
21 as an LBA) and an amount of data (such as a number of LBA quanta, e.g., sectors) to read; in
22 response the SSD provides read status and/or read data. For another example, the commands
23 include a write command specifying an address (such as an LBA) and an amount of data (such
24 as a number of LBA quanta, e.g., sectors) to write; in response the SSD provides write status
25 and/or requests write data and optionally subsequently provides write status. For yet another
26 example, the commands include a de-allocation command (e.g. a trim command) specifying one
27 or more addresses (such as one or more LBAs) that no longer need be allocated; in response the
28 SSD modifies the Map accordingly and optionally provides de-allocation status. In some
29 contexts an ATA compatible TRIM command is an exemplary de-allocation command. For yet
30 another example, the commands include a super capacitor test command or a data hardening
31 success query; in response, the SSD provides appropriate status. In some embodiments, Host
32 Interfaces **111** is compatible with a SATA protocol and, using NCQ commands, is enabled to
33 have up to 32 pending commands, each with a unique tag represented as a number from 0 to 31.
34 In some embodiments, Tag Tracking **113** is enabled to associate an external tag for a command
35 received via External Interfaces **110** with an internal tag used to track the command during
36 processing by SSD Controller **100**.

1
2 [0092] According to various embodiments, one or more of: Data Processing **121**
3 optionally and/or selectively processes some or all data sent between Buffer **131** and External
4 Interfaces **110**; and Data Processing **121** optionally and/or selectively processes data stored in
5 Buffer **131**. In some embodiments, Data Processing **121** uses one or more Engines **123** to
6 perform one or more of: formatting; reformatting; transcoding; and any other data processing
7 and/or manipulation task.

8
9 [0093] Buffer **131** stores data sent to/from External Interfaces **110** from/to Device
10 Interfaces **190**. In some embodiments, Buffer **131** additionally stores system data, such as some
11 or all map tables, used by SSD Controller **100** to manage one or more of Flash Device **192**. In
12 various embodiments, Buffer **131** has one or more of: Memory **137** used for temporary storage
13 of data; DMA **133** used to control movement of data to and/or from Buffer **131**; and ECC-X **135**
14 used to provide higher-level error correction and/or redundancy functions; and other data
15 movement and/or manipulation functions. An example of a higher-level redundancy function is
16 a RAID-like capability, where redundancy is at a flash device (e.g., multiple ones of Flash
17 Device **192**) level and/or a flash die (e.g., Flash Die **194**) level instead of at a disk level.

18
19 [0094] According to various embodiments, one or more of: ECC **161** optionally and/or
20 selectively processes some or all data sent between Buffer **131** and Device Interfaces **190**; and
21 ECC **161** optionally and/or selectively processes data stored in Buffer **131**. In some
22 embodiments, ECC **161** is used to provide lower-level error correction and/or redundancy
23 functions, such as in accordance with one or more ECC techniques. In some embodiments, ECC
24 **161** implements one or more of: a CRC code; a Hamming code; an RS code; a BCH code; an
25 LDPC code; a Viterbi code; a trellis code; a hard-decision code; a soft-decision code; an erasure-
26 based code; any error detecting and/or correcting code; and any combination of the preceding.
27 In some embodiments, ECC **161** includes one or more decoders (such as LDPC decoders).

28
29 [0095] Device Interface Logic **191** controls instances of Flash Device **192** via Device
30 Interfaces **190**. Device Interface Logic **191** is enabled to send data to/from the instances of Flash
31 Device **192** according to a protocol of Flash Device **192**. Device Interface Logic **191** includes
32 Scheduling **193** to selectively sequence control of the instances of Flash Device **192** via Device
33 Interfaces **190**. For example, in some embodiments, Scheduling **193** is enabled to queue
34 operations to the instances of Flash Device **192**, and to selectively send the operations to
35 individual ones of the instances of Flash Device **192** (or Flash Die **194**) as individual ones of the
36 instances of Flash Device **192** (or Flash Die **194**) are available.

1
2 [0096] Map 141 converts between data addressing used on External Interfaces 110 and
3 data addressing used on Device Interfaces 190, using Table 143 to map external data addresses
4 to locations in NVM 199. For example, in some embodiments, Map 141 converts LBAs used on
5 External Interfaces 110 to block and/or page addresses targeting one or more Flash Die 194, via
6 mapping provided by Table 143. For LBAs that have never been written since drive
7 manufacture or de-allocation, the Map points to a default value to return if the LBAs are read.
8 For example, when processing a de-allocation command, the Map is modified so that entries
9 corresponding to the de-allocated LBAs point to one of the default values. In various
10 embodiments, there are various default values, each having a corresponding pointer. The
11 plurality of default values enables reading some de-allocated LBAs (such as in a first range) as
12 one default value, while reading other de-allocated LBAs (such as in a second range) as another
13 default value. The default values, in various embodiments, are defined by flash memory,
14 hardware, firmware, command and/or primitive arguments and/or parameters, programmable
15 registers, or various combinations thereof.

16

17 [0097] In some embodiments, Map 141 uses Table 143 to perform and/or to look up
18 translations between addresses used on External Interfaces 110 and data addressing used on
19 Device Interfaces 190. According to various embodiments, Table 143 is one or more of: a one-
20 level map; a two-level map; a multi-level map; a map cache; a compressed map; any type of
21 mapping from one address space to another; and any combination of the foregoing. According
22 to various embodiments, Table 143 includes one or more of: static random access memory;
23 dynamic random access memory; NVM (such as flash memory); cache memory; on-chip
24 memory; off-chip memory; and any combination of the foregoing.

25

26 [0098] In some embodiments, Recycler 151 performs garbage collection. For example,
27 in some embodiments, instances of Flash Device 192 contain blocks that must be erased before
28 the blocks are re-writeable. Recycler 151 is enabled to determine which portions of the
29 instances of Flash Device 192 are actively in use (e.g., allocated instead of de-allocated), such as
30 by scanning a map maintained by Map 141, and to make unused (e.g., de-allocated) portions of
31 the instances of Flash Device 192 available for writing by erasing them. In further
32 embodiments, Recycler 151 is enabled to move data stored within instances of Flash Device 192
33 to make larger contiguous portions of the instances of Flash Device 192 available for writing.

34

35 [0099] In some embodiments, instances of Flash Device 192 are selectively and/or
36 dynamically configured, managed, and/or used to have one or more bands for storing data of

1 different types and/or properties. A number, arrangement, size, and type of the bands are
2 dynamically changeable. For example, data from a computing host is written into a hot (active)
3 band, while data from Recycler 151 is written into a cold (less active) band. In some usage
4 scenarios, if the computing host writes a long, sequential stream, then a size of the hot band
5 grows, whereas if the computing host does random writes or few writes, then a size of the cold
6 band grows.

7

8 [0100] CPU 171 controls various portions of SSD Controller 100. CPU 171 includes
9 CPU Core 172. CPU Core 172 is, according to various embodiments, one or more single-core or
10 multi-core processors. The individual processors cores in CPU Core 172 are, in some
11 embodiments, multi-threaded. CPU Core 172 includes instruction and/or data caches and/or
12 memories. For example, the instruction memory contains instructions to enable CPU Core 172
13 to execute programs (e.g. software sometimes called firmware) to control SSD Controller 100.
14 In some embodiments, some or all of the firmware executed by CPU Core 172 is stored on
15 instances of Flash Device 192 (as illustrated, e.g., as Firmware 106 of NVM 199 in Fig. 1B).

16

17 [0101] In various embodiments, CPU 171 further includes: Command Management
18 173 to track and control commands received via External Interfaces 110 while the commands are
19 in progress; Buffer Management 175 to control allocation and use of Buffer 131; Translation
20 Management 177 to control Map 141; Coherency Management 179 to control consistency of
21 data addressing and to avoid conflicts such as between external data accesses and recycle data
22 accesses; Device Management 181 to control Device Interface Logic 191; Identity Management
23 182 to control modification and communication of identify information, and optionally other
24 management units. None, any, or all of the management functions performed by CPU 171 are,
25 according to various embodiments, controlled and/or managed by hardware, by software (such
26 as firmware executing on CPU Core 172 or on a host connected via External Interfaces 110), or
27 any combination thereof.

28

29 [0102] In some embodiments, CPU 171 is enabled to perform other management tasks,
30 such as one or more of: gathering and/or reporting performance statistics; implementing
31 SMART; controlling power sequencing, controlling and/or monitoring and/or adjusting power
32 consumption; responding to power failures; controlling and/or monitoring and/or adjusting clock
33 rates; and other management tasks.

34

35 [0103] Various embodiments include a computing-host flash memory controller that is
36 similar to SSD Controller 100 and is compatible with operation with various computing hosts,

1 such as via adaptation of Host Interfaces **111** and/or External Interfaces **110**. The various
2 computing hosts include one or any combination of a computer, a workstation computer, a
3 server computer, a storage server, a SAN, a NAS device, a DAS device, a storage appliance, a
4 PC, a laptop computer, a notebook computer, a netbook computer, a tablet device or computer,
5 an ultrabook computer, an electronic reading device (such as an e-reader), a PDA, a navigation
6 system, a (handheld) GPS device, an automotive control system, an automotive media control
7 system or computer, a printer, copier or fax machine or all-in-one device, a POS device, a cash-
8 register, a media player, a television, a media recorder, a DVR, a digital camera, a cellular
9 handset, a cordless telephone handset, and an electronic game.

10

11 [0104] In various embodiments, all or any portions of an SSD controller (or a
12 computing-host flash memory controller) are implemented on a single IC, a single die of a multi-
13 die IC, a plurality of dice of a multi-die IC, or a plurality of ICs. For example, Buffer **131** is
14 implemented on a same die as other elements of SSD Controller **100**. For another example,
15 Buffer **131** is implemented on a different die than other elements of SSD Controller **100**.

16

17 [0105] Fig. 1B illustrates selected details of various embodiments of systems including
18 one or more instances of the SSD of Fig. 1A. SSD **101** includes SSD Controller **100** coupled to
19 NVM **199** via Device Interfaces **190**. The figure illustrates various classes of embodiments: a
20 single SSD coupled directly to a host, a plurality of SSDs each respectively coupled directly to a
21 host via respective external interfaces, and one or more SSDs coupled indirectly to a host via
22 various interconnection elements.

23

24 [0106] As an example embodiment of a single SSD coupled directly to a host, one
25 instance of SSD **101** is coupled directly to Host **102** via External Interfaces **110** (e.g. Switch /
26 Fabric / Intermediate Controller **103** is omitted, bypassed, or passed-through). As an example
27 embodiment of a plurality of SSDs each coupled directly to a host via respective external
28 interfaces, each of a plurality of instances of SSD **101** is respectively coupled directly to Host
29 **102** via a respective instance of External Interfaces **110** (e.g. Switch / Fabric / Intermediate
30 Controller **103** is omitted, bypassed, or passed-through). As an example embodiment of one or
31 more SSDs coupled indirectly to a host via various interconnection elements, each of one or
32 more instances of SSD **101** is respectively coupled indirectly to Host **102**. Each indirect
33 coupling is via a respective instance of External Interfaces **110** coupled to Switch / Fabric /
34 Intermediate Controller **103**, and Intermediate Interfaces **104** coupling to Host **102**.

35

1 [0107] Some of the embodiments including Switch / Fabric / Intermediate Controller
2 103 also include Card Memory 112C coupled via Memory Interface 180 and accessible by the
3 SSDs. In various embodiments, one or more of the SSDs, the Switch / Fabric / Intermediate
4 Controller, and/or the Card Memory are included on a physically identifiable module, card, or
5 pluggable element (e.g. I/O Card 116). In some embodiments, SSD 101 (or variations thereof)
6 corresponds to a SAS drive or a SATA drive that is coupled to an initiator operating as Host 102.
7

8 [0108] Host 102 is enabled to execute various elements of Host Software 115, such as
9 various combinations of OS 105, Driver 107, Application 109, and Multi-Device Management
10 Software 114. Dotted-arrow 107D is representative of Host Software ↔ I/O Device
11 Communication, e.g. data sent/received to/from one or more of the instances of SSD 101 and
12 from/to any one or more of OS 105 via Driver 107, Driver 107, and Application 109, either via
13 Driver 107, or directly as a VF.
14

15 [0109] OS 105 includes and/or is enabled to operate with drivers (illustrated
16 conceptually by Driver 107) for interfacing with the SSD. Various versions of Windows (e.g.
17 95, 98, ME, NT, XP, 2000, Server, Vista, and 7), various versions of Linux (e.g. Red Hat,
18 Debian, and Ubuntu), and various versions of MacOS (e.g. 8, 9 and X) are examples of OS 105.
19 In various embodiments, the drivers are standard and/or generic drivers (sometimes termed
20 “shrink-wrapped” or “pre-installed”) operable with a standard interface and/or protocol such as
21 SATA, AHCI, or NVM Express, or are optionally customized and/or vendor specific to enable
22 use of commands specific to SSD 101. Some drives and/or drivers have pass-through modes to
23 enable application-level programs, such as Application 109 via Optimized NAND Access
24 (sometimes termed ONA) or Direct NAND Access (sometimes termed DNA) techniques, to
25 communicate commands directly to SSD 101, enabling a customized application to use
26 commands specific to SSD 101 even with a generic driver. ONA techniques include one or
27 more of: use of non-standard modifiers (hints); use of vendor-specific commands;
28 communication of non-standard statistics, such as actual NVM usage according to
29 compressibility; and other techniques. DNA techniques include one or more of: use of non-
30 standard commands or vendor-specific providing unmapped read, write, and/or erase access to
31 the NVM; use of non-standard or vendor-specific commands providing more direct access to the
32 NVM, such as by bypassing formatting of data that the I/O device would otherwise do; and other
33 techniques. Examples of the driver are a driver without ONA or DNA support, an ONA-enabled
34 driver, a DNA-enabled driver, and an ONA/DNA-enabled driver. Further examples of the driver
35 are a vendor-provided, vendor-developed, and/or vendor-enhanced driver, and a client-provided,
36 client-developed, and/or client-enhanced driver.

1
2 [0110] Examples of the application-level programs are an application without ONA or
3 DNA support, an ONA-enabled application, a DNA-enabled application, and an ONA/DNA-
4 enabled application. Dotted-arrow **109D** is representative of Application \leftrightarrow I/O Device
5 Communication (e.g. bypass via a driver or bypass via a VF for an application), e.g. an ONA-
6 enabled application and an ONA-enabled driver communicating with an SSD, such as without
7 the application using the OS as an intermediary. Dotted-arrow **109V** is representative of
8 Application \leftrightarrow I/O Device Communication (e.g. bypass via a VF for an application), e.g. a
9 DNA-enabled application and a DNA-enabled driver communicating with an SSD, such as
10 without the application using the OS or the driver as intermediaries.

11

12 [0111] One or more portions of NVM **199** are used, in some embodiments, for
13 firmware storage, e.g. Firmware **106**. The firmware storage includes one or more firmware
14 images (or portions thereof). A firmware image has, for example, one or more images of
15 firmware, executed, e.g., by CPU Core **172** of SSD Controller **100**. A firmware image has, for
16 another example, one or more images of constants, parameter values, and NVM device
17 information, referenced, e.g. by the CPU core during the firmware execution. The images of
18 firmware correspond, e.g., to a current firmware image and zero or more previous (with respect
19 to firmware updates) firmware images. In various embodiments, the firmware provides for
20 generic, standard, ONA, and/or DNA operating modes. In some embodiments, one or more of
21 the firmware operating modes are enabled (e.g. one or more APIs are “unlocked”) via keys or
22 various software techniques, optionally communicated and/or provided by a driver.

23

24 [0112] In some embodiments lacking the Switch / Fabric / Intermediate Controller, the
25 SSD is coupled to the Host directly via External Interfaces **110**. In various embodiments, SSD
26 Controller **100** is coupled to the Host via one or more intermediate levels of other controllers,
27 such as a RAID controller. In some embodiments, SSD **101** (or variations thereof) corresponds
28 to a SAS drive or a SATA drive and Switch / Fabric / Intermediate Controller **103** corresponds
29 to an expander that is in turn coupled to an initiator, or alternatively Switch / Fabric /
30 Intermediate Controller **103** corresponds to a bridge that is indirectly coupled to an initiator via
31 an expander. In some embodiments, Switch / Fabric / Intermediate Controller **103** includes one
32 or more PCIe switches and/or fabrics.

33

34 [0113] In various embodiments, such as some of the embodiments where Host **102** is a
35 computing host (e.g. a computer, a workstation computer, a server computer, a storage server, a
36 SAN, a NAS device, a DAS device, a storage appliance, a PC, a laptop computer, a notebook

1 computer, and/or a netbook computer), the computing host is optionally enabled to communicate
2 (e.g. via optional I/O & Storage Devices/Resources 117 and optional LAN/WAN 119) with one
3 or more local and/or remote servers (e.g. optional Servers 118). The communication enables, for
4 example, local and/or remote access, management, and/or usage of any one or more of SSD 101
5 elements. In some embodiments, the communication is wholly or partially via Ethernet. In
6 some embodiments, the communication is wholly or partially via Fibre Channel. LAN/WAN
7 119 is representative, in various embodiments, of one or more Local and/or Wide Area
8 Networks, such as any one or more of a network in a server farm, a network coupling server
9 farms, a metro-area network, and the Internet.

10

11 [0114] In various embodiments, an SSD controller and/or a computing-host flash
12 memory controller in combination with one or more NVMs are implemented as a non-volatile
13 storage component, such as a USB storage component, a CF storage component, an MMC
14 storage component, an eMMC storage component, a Thunderbolt storage component, a UFS
15 storage component, an SD storage component, a Memory Stick storage component, and an xD-
16 picture card storage component.

17

18 [0115] In various embodiments, all or any portions of an SSD controller (or a
19 computing-host flash memory controller), or functions thereof, are implemented in a host that
20 the controller is to be coupled with (e.g., Host 102 of Fig. 1B). In various embodiments, all or
21 any portions of an SSD controller (or a computing-host flash memory controller), or functions
22 thereof, are implemented via hardware (e.g., logic circuitry), software and/or firmware (e.g.,
23 driver software or SSD control firmware), or any combination thereof. For example,
24 functionality of or associated with an ECC unit (such as similar to ECC 161 and/or ECC-X 135
25 of Fig. 1A) is implemented partially via software on a host and partially via a combination of
26 firmware and hardware in an SSD controller. For another example, functionality of or
27 associated with a recycler unit (such as similar to Recycler 151 of Fig. 1A) is implemented
28 partially via software on a host and partially via hardware in a computing-host flash memory
29 controller.

30

31

32 VARIABLE OVER-PROVISIONING (OP)

33

34 [0116] Fig. 2 illustrates selected details of various host and system allocations of Flash
35 Memory 200, including for OP use, according to various embodiments of a system using
36 variable OP for managing NVMs. The host allocation is used to store host data, such as data

1 received/provided from/to a host coupled to an SSD using flash memory for non-volatile
2 storage. A system allocation (e.g. System Allocation **201**) is used to store system data, such as
3 relating to management or operation of an SSD controller (e.g. periodic copies of all or portions
4 of information in map **141** of Fig. 1A), an SSD, a flash memory, or any portions thereof.

5

6 [0117] The figure illustrates Flash Memory **200** used according to three allocation
7 scenarios (Initial, Same, and Increased Allocations **210A**, **210B**, and **210C**, respectively). Each
8 of the allocation scenarios illustrates, at a highest level of detail, three elements to the allocation:
9 host, system, and OP. In Initial Allocation **210A** and Same Allocation **210B**, the host, system,
10 and OP allocations are illustrated respectively as Host Allocation **202A**, System Allocation **201**,
11 and OP (system + host) Allocation **203A**. In Increased Allocation **210C**, the host, system, and
12 OP allocations are illustrated respectively as Host Allocation **202C**, System Allocation **201**, and
13 OP (system + host) Allocation **203C**. At a further level of detail, the various allocations differ
14 according to allocating between host OP allocation and system OP allocation, illustrated as Host
15 OP Allocations **205A** and **205B**, and System OP Allocations **204A**, **204B**, and **204C**, as
16 described following.

17

18 [0118] In operation, allocation begins according to a starting state (e.g. Initial
19 Allocation **210A**). In response to an event, an agent changes the allocation to a different state
20 (e.g. Same Allocation **210B** or Increased Allocation **210C**). Various embodiments are according
21 to various operating modes. In a first operating mode, host allocation remains unchanged and
22 OP (system + host) allocation remains unchanged, but is redistributed between system and host
23 OP allocations. In a second operating mode, host allocation is changed to enable corresponding
24 changes to OP (system + host) allocation that is redistributed between system and host OP
25 allocations.

26

27 [0119] In the first operating mode, after allocation has begun according to Initial
28 Allocation **210A**, allocation is then dynamically altered to Same Allocation **210B**. At the
29 highest level of detail, the three elements of the allocation remain unchanged. Specifically the
30 host, system, and OP allocations remain unchanged at, respectively, Host Allocation **202A**,
31 System Allocation **201**, and OP (system + host) Allocation **203A**. However, at the further level
32 of detail, the OP allocation is internally altered by decreasing the host OP allocation by an
33 amount and increasing the system OP allocation by the amount. Specifically the host OP
34 allocation decreases from Host OP Allocation **205A** to Host OP Allocation **205B** by amount
35 System OP Allocation Delta **208**, and the system OP allocation increases from System OP
36 Allocation **204A** to System OP Allocation **204B** by amount System OP Allocation Delta **208**.

1
2 [0120] After further operation, allocation is then dynamically altered to return to Initial
3 Allocation **210A**, including the host OP allocation being Host OP Allocation **205A** and the
4 system OP allocation being System OP Allocation **204A**. Other operating scenarios applicable
5 to the first operating mode include any change in allocation of the OP allocation between host
6 OP and system OP uses (while leaving the host allocation and the system allocation unchanged),
7 without restriction to any of the particular allocations illustrated in the figure.

8
9 [0121] In the second operating mode, after allocation has begun according to Initial
10 Allocation **210A**, allocation is then dynamically altered to Increased Allocation **210C**. At the
11 highest level of detail, the three elements of the allocation are altered so that the host allocation
12 is decreased by an amount, and the OP allocation is increased by the amount. At the further
13 level of detail, the host OP allocation remains unchanged and the system OP allocation increases
14 by the amount. Specifically the host allocation decreases from Host Allocation **202A** to Host
15 Allocation **202C** by amount System OP Allocation Delta **209**, and the OP allocation increases
16 from OP Allocation **203A** to OP Allocation **203C** by amount System OP Allocation Delta **209**.
17 At the further level of detail, the host OP allocation remains unchanged at Host OP Allocation
18 **205A**, and the system OP allocation increases from System OP Allocation **204A** to System OP
19 Allocation **204C** by amount System OP Allocation Delta **209**.

20
21 [0122] After further operation, allocation is then dynamically altered to return to Initial
22 Allocation **210A**, including the host OP allocation being Host OP Allocation **205A** and the
23 system OP allocation being System OP Allocation **204A**. Other operating scenarios applicable
24 to the second operating mode include any change in allocation of the OP allocation between host
25 OP and system OP uses (in conjunction with increasing/decreasing the host allocation), without
26 restriction to any of the particular allocations illustrated in the figure.

27
28 [0123] In a third operating mode (not illustrated), instead of (and/or in addition to)
29 changing the host allocation to accommodate a change in the OP allocation (as in the second
30 operating mode), the system allocation is changed. Thus a decrease/increase in system
31 allocation enables an increase/decrease in system OP allocation, or alternatively a
32 decrease/increase in host allocation in combination with a decrease/increase in system allocation
33 enables an increase/decrease in host and/or system OP allocation. Other operating modes that
34 are any combination of the first through the third operating modes are contemplated, included
35 operating modes that dynamically switch between any of the first through the third operating
36 modes, according to various operating criteria and/or characteristics.

1
2 [0124] There are several events that lead to an allocation state change and
3 corresponding agents that effect the allocation state change. For example, an SSD controller
4 (e.g. SSD controller 100 of Fig. 1A) determines that an increased (or decreased) amount of flash
5 memory (e.g. NVM 199 of Fig. 1A) is to be used to store host data, such as due to increased (or
6 decreased) usage via changed compression, deduplication, or transformation. In response, the
7 SSD controller increases (or decreases) system OP allocation (e.g. such as from System OP
8 Allocation 204A to System OP Allocation 204C of Fig. 2). For another example, a processor
9 (e.g. CPU 171 of Fig. 1A) within a flash memory based storage sub-system (e.g. SSD 101 of
10 Fig. 1B) receives a command via a storage interface (e.g. external interfaces 110 of Fig. 1B) of
11 the storage sub-system. The command (e.g. an ATA compatible TRIM command) specifies that
12 particular portions of the storage sub-system are unused (e.g. free), and that any data stored
13 therein is no longer needed. In response, the processor decreases allocation of the flash memory
14 to host data and/or host OP and increases allocation of the flash memory to system OP. In
15 various embodiments, the allocation decrease to host data and/or host OP (and optionally and/or
16 selectively an increase to system OP) is postponed until the unused portions have been garbage
17 collected, recycled, and/or erased.

18
19 [0125] For yet another example, a monitoring sub-system within an SSD controller
20 (e.g. all or any portions of scheduling 193 of Fig. 1A) determines a current ratio of bandwidths
21 of system writes to host writes has changed with respect to a previous ratio. In response, the
22 monitoring sub-system requests that an allocation agent (e.g. all or any portions of recycler 151
23 of Fig. 1A) of the SSD controller alter allocation between system OP and host OP in accordance
24 with the current ratio. E.g. if the current ratio has increased (decreased), then the allocation
25 agent is requested to increase (decrease) the system OP and/or to decrease (increase) the host
26 OP. In various embodiments the increase (decrease) is a linear function of the current ratio, a
27 linear function of a ratio of the current and the previous ratios, an inverse function of the current
28 ratio, an inverse function of the current and the previous ratios, a non-linear function of the
29 current ratio and/or the ratio of the current and the previous ratios, any combination thereof, or
30 any approximation(s) thereof.

31
32
33 DYNAMICALLY VARIABLE OVER-PROVISIONING (OP) AND DATA ENTROPY

34
35 [0126] Fig. 3A and Fig. 3B illustrate selected details of various embodiments of
36 dynamically varying host and system OPs as relating to dynamically varying data entropy. The

1 host OP and system OP dynamically vary as one or more host-allocation functions of
2 (dynamically varying) host allocation as well as one or more data-rate-allocation functions of
3 (dynamically varying) data rates (e.g. of host data and system data). The host allocation
4 variation is illustrated as linear with respect to data entropy, for convenience of explanation.
5

6 [0127] For instance, if host allocation decreases/increases by an amount, then the
7 amount is allocated to/from combined OP, and a resultant combined OP is dynamically allocated
8 between host OP and system OP according to one or more data-rate-allocation functions of one
9 or more data rates. Examples of the data-rate-allocation functions are a linear function, an
10 inverse function, a non-linear function, or any combinations thereof. Examples of the data rates
11 are a host data rate, a system data rate, bandwidth of host data writes to flash memory,
12 bandwidth of system data writes to flash memory, bandwidth of total writes to flash memory, or
13 any combinations thereof.

14

15 [0128] As illustrated in Fig. 3A and Fig. 3B, a (combined) OP allocation dynamically
16 varies due to a corresponding host allocation dynamic variation. The dynamically varying
17 (combined) OP allocation is then dynamically variably allocated between system OP and host
18 OP, e.g., according to a ratio of bandwidth of system data writes to flash memory and bandwidth
19 of host data writes to flash memory. Therefore, for a particular fractional decrease/increase of
20 host allocation (and at a same ratio), a system OP allocation increases/decreases in accordance
21 with the particular fraction. In various embodiments, the ratio is optionally scaled and/or an
22 offset is included in the allocation between host OP and system OP. In some embodiments, the
23 allocations are further in accordance with one or more of respective minimum/maximum values
24 and/or respective allocation quanta.

25

26 [0129] Common to both figures, the horizontal axis represents Data Entropy 320,
27 increasing left (0%) to right (100%). Lower data entropy corresponds to host data that is, e.g.,
28 relatively highly compressible, relatively highly compactable via deduplication, and/or relatively
29 less random. Higher data entropy corresponds to host data that is, e.g. relatively highly
30 incompressible, relatively lowly duplicative (and thus not highly compactable via
31 deduplication), and/or relatively more random. Two particular data entropy values are
32 illustrated as G% 330G and H% 330H.

33

34 [0130] In various embodiments, data entropy is conceptual, e.g. no explicit
35 measurements of data entropy are made. Instead, data entropy is a representation of how host
36 data size varies with varying compression, deduplication, or other transformations that alter

1 amount of flash memory used to store host data. For example, flash memory used to store host
2 data increases (decreases) as data entropy increases (decreases). For instance, when illustrating
3 a relationship between data entropy and host data size (e.g. as a host data allocation), a scale for
4 data entropy is interpreted as linear, logarithmic, square-law, arbitrarily variable, or any
5 combinations thereof.

6

7 [0131] The vertical axis represents Physical Size **310**, increasing from bottom (0%) to
8 top (100%). Length along the vertical axis corresponds to flash memory used and/or allocated to
9 be used; e.g. a shorter/longer length corresponds to less/more flash memory used to store a given
10 type of information (e.g. host data or system data). In some embodiments, a particular length
11 corresponds to a number of units of flash memory allocated to a particular usage (e.g. to host
12 data, system data, or OP).

13

14 [0132] A host allocation function that varies with respect to data entropy is illustrated
15 as Host Allocation **302V**. As data entropy increases, physical size of flash memory used to store
16 corresponding host data (e.g. linearly) increases, and vice-versa. An invariant system allocation
17 function that remains fixed as data entropy increases (decreases) is illustrated as System
18 Allocation **301**. Embodiments (not illustrated) are contemplated where system allocation varies
19 according to data entropy, such as increasing (decreasing) as data entropy (or any one or more
20 metrics associated with system data size and/or host data size) increases (decreases).

21

22 [0133] Fig. 3A exemplifies two “operating points” (one corresponding to G% **330G**
23 and another to H% **330H**) of a first dynamically varying allocation of System OP Allocation
24 (variable) **303V1** corresponding to a linear function of a first value of a ratio. Fig. 3B
25 exemplifies two “operating points” (one corresponding to G% **330G** and another to H% **330H**)
26 of a second dynamically varying allocation of System OP Allocation (variable) **303V2**
27 corresponding to a linear function of a second value of the ratio. For comparison, Fig. 3A
28 includes a dashed-line of System OP Allocation (variable) **303V2** (having two operating points
29 detailed in Fig. 3B), and Fig. 3B includes a dashed-line of System OP Allocation (variable)
30 **303V1** (having two operating points detailed in Fig. 3A).

31

32 [0134] In various embodiments, the ratio is a ratio of a system rate to a host rate. In
33 various embodiments, the system and/or the host rate is a function of a current and/or previous
34 data rate and/or write bandwidth. For example, the ratio is system write data bandwidth
35 averaged over a time interval divided by host write data bandwidth averaged over the time

1 interval. For another example, the ratio is system data rate at a previous point in time divided by
2 host data rate at the previous point in time.

3

4 [0135] In Fig. 3A, the first operating point (corresponding to G% **330G**) is with respect
5 to an intersection of G% data entropy with Host Allocation **302V** and System OP Allocation
6 (variable) **303V1**. The first operating point includes Host Allocation **302G1** and (combined) OP
7 Allocation **303G1** allocated between Host OP Allocation **305G1** and System OP Allocation
8 **304G1**. The second operating point (corresponding to H% **330H**) is with respect to an
9 intersection of H% data entropy with Host Allocation **302V** and System OP Allocation
10 (variable) **303V1**. The second operating point includes Host Allocation **302H1** and (combined)
11 OP Allocation **303H1** allocated between Host OP Allocation **305H1** and System OP Allocation
12 **304H1**.

13

14 [0136] Comparing the first and the second operating points, as data entropy increases
15 from G% to H%, flash memory used to store host data increases from Host Allocation **302G1** to
16 Host Allocation **302H1**. As less flash memory is then available for OP usage, in response
17 (combined) OP allocation decreases from (combined) OP Allocation **303G1** to (combined) OP
18 Allocation **303H1**, and corresponding allocations of host OP and system OP are decreased.
19 Specifically, Host OP Allocation **305G1** is decreased to Host OP Allocation **305H1**, and System
20 OP Allocation **304G1** is decreased to System OP Allocation **304H1**. The decreases to host OP
21 and system OP allocation are in accordance with an allocation of OP resources between host OP
22 and system OP in accordance with a first value of a ratio, such as a ratio of system write
23 bandwidth to host data bandwidth. The allocation is conceptually represented by System OP
24 Allocation (variable) **303V1**, illustrated as a linear function of Host Allocation **302V**.

25

26 [0137] Fig. 3B illustrates techniques similar to techniques illustrated by Fig. 3A. The
27 first and second operating points are with respect to respective intersections of G% and H% data
28 entropies with Host Allocation **302V** and System OP Allocation (variable) **303V2**. The first
29 operating point includes Host Allocation **302G2** and (combined) OP Allocation **303G2** allocated
30 between Host OP Allocation **305G2** and System OP Allocation **304G2**. The second operating
31 point includes Host Allocation **302H2** and (combined) OP Allocation **303H2** allocated between
32 Host OP Allocation **305H2** and System OP Allocation **304H2**.

33

34 [0138] Similar to Fig. 3A, as illustrated in Fig. 3B, an increase in data entropy (e.g.
35 from G% to H%) results in an increase in host allocation (e.g. from Host Allocation **302G2** to
36 Host Allocation **302H2**), in turn resulting in a decrease to (combined) OP allocation (e.g. from

1 (combined) OP Allocation **303G2** to (combined) OP Allocation **303H2**) that is then reallocated
2 between system OP allocations (e.g. from System OP Allocation **304G2** to System OP
3 Allocation **304H2**) and host OP allocations (e.g. from Host OP Allocation **305G2** to Host OP
4 Allocation **305H2**). The decreases to system OP and host OP allocation are in accordance with
5 an allocation of OP resources between system OP and host OP in accordance with the second
6 value of the ratio associated with Fig. 3A.

7

8 [0139] Comparing Fig. 3A to Fig. 3B, the slope of System OP Allocation (variable)
9 **303V1** is less than that of System OP Allocation (variable) **303V2**, corresponding to the first
10 value of the ratio being less than the second value of the ratio. Consider embodiments where the
11 ratio is a system data rate divided by a host data rate. In a usage scenario where the system data
12 rate dynamically varies, while the host data rate is unchanging, System OP Allocation (variable)
13 **303V1** corresponds to a lower system data rate than System OP Allocation (variable) **303V2**.
14 Thus as the system data rate dynamically increases (with respect to a constant host data rate), the
15 system OP allocation increases from System OP Allocation **304G1** to System OP Allocation
16 **304G2** (or alternatively from System OP Allocation **304H1** to System OP Allocation **304H2**).
17 Correspondingly, the host OP allocation decreases from Host OP Allocation **305G1** to Host OP
18 Allocation **305G2** (or alternatively from Host OP Allocation **305H1** to Host OP Allocation
19 **305H2**). Alternatively, in a usage scenario where the host data rate dynamically varies, while
20 the system data rate is unchanging, the foregoing system and host OP allocation changes
21 likewise occur (System OP Allocation (variable) **303V1** corresponds to a higher host data rate
22 than System OP Allocation (variable) **303V2**).
23

24 [0140] The first and the second operating points (of Fig. 3A and Fig. 3B), as well as the
25 foregoing description of data entropy increasing from G% to H%, are examples only. For
26 another example, in some usage scenarios, data entropy decreases from H% to G%, and flash
27 memory allocation is altered from being in accordance with the second operating point (H%) to
28 being in accordance with the first operating point (G%), such that an allocation of flash memory
29 to system OP usage increases. Further, a multiplicity of possible operating points along the data
30 entropy axis are possible (not illustrated), limited only by implementation details. Still further,
31 in some circumstances, operational scenarios dynamically switch among various operational
32 scenarios and corresponding operating points dynamically, based, e.g., on various operating
33 characteristics of an SSD and/or an SSD controller.
34

35 [0141] For comparison and reference, Figs. 3A and 3B identically illustrate System
36 Allocation (fixed) **301F** that is fixed (e.g. invariant with respect to data entropy). Some

1 embodiments and/or usage scenarios include an operating mode and/or operating sub-mode,
2 where system OP resources are fixed (e.g. invariant with respect to data entropy), as illustrated
3 conceptually by System OP Allocation (fixed) **304F**. In some embodiments, System OP
4 Allocation (fixed) **304F** is representative of all or a portion of a minimum system OP allocation.
5

6 [0142] Fig. 3A and Fig. 3B illustrate aspects of several examples of dynamically
7 varying system OP (e.g. via dynamic allocation of OP resources between host OP and system
8 OP) as a result of dynamic variance of host allocation, that is in turn affected by dynamic
9 variance in data entropy. Some of the aspects are conceptually represented with respect to either
10 of the figures alone, and some of the aspects are conceptually represented with both figures in
11 combination.

12

13 [0143] With respect to Fig. 3A alone (or Fig. 3B alone), the dynamically varying
14 system OP is conceptually represented with respect to dynamic changes between various host
15 allocations (corresponding to various data entropy operating points), while a ratio of system data
16 rate to host data rate remains fixed. Consider beginning operating at a host allocation
17 corresponding to G% entropy, such as the intersection with Host Allocation **302V** and a system
18 OP allocation function represented by System OP Allocation (variable) **303V1**. Then continue
19 by transitioning to operating at a host allocation corresponding to H% entropy, while continuing
20 to allocate and/or allocate system OP resources according to the function represented by System
21 OP Allocation (variable) **303V1**. Then continue by transitioning back to operating at the host
22 allocation corresponding to G% entropy, while continuing to allocate and/or allocate system OP
23 resources according to the function represented by System OP Allocation (variable) **303V1**. The
24 system OP is dynamically varied from System OP Allocation **304G1** to System OP Allocation
25 **304H1** and then back to System OP Allocation **304G1**, in response to changes in host allocation
26 that are in turn related to changes in data entropy operating points.

27

28 [0144] With respect to Fig. 3A in combination with Fig. 3B, the dynamically varying
29 system OP is conceptually represented with respect to dynamic changes between allocation of
30 OP resources between host OP and system OP uses based on differences and/or changes in
31 system and/or host data rates and/or write bandwidths, one or more functions used to determine
32 the allocation, or any combination thereof. Consider beginning operating at a context as
33 illustrated in Fig. 3A of G% data entropy and System OP Allocation (variable) **303V1**. Then
34 continue by transitioning to operating at a context as illustrated by Fig. 3B of G% data entropy
35 and System OP Allocation (variable) **303V2**. Then continue by transitioning back to operating
36 at the context as illustrated Fig. 3A of G% data entropy and System OP Allocation (variable)

1 **303V1.** The system OP is dynamically varied from System OP Allocation **304G1** to System OP
2 Allocation **304G2** and then back to System OP Allocation **304G1**, in response to changes in,
3 e.g., system data write bandwidth, host data write bandwidth, and/or ratio(s) thereof.

4

5 **[0145]** With respect to Fig. 3A in combination with Fig. 3B, the dynamically varying
6 system OP is also conceptually represented with respect to dynamic changes between various
7 host allocations (corresponding to various data entropy operating points) in combination with
8 dynamic changes between allocation of OP resources between host OP and system OP uses
9 based on differences and/or changes in system and/or host data rates and/or write bandwidths
10 and/or ratios thereof. Consider beginning operating at a context as illustrated in Fig. 3A of G%
11 data entropy and System OP Allocation (variable) **303V1**. Then continue by transitioning to
12 operating at a context as illustrated by Fig. 3B of H% data entropy and System OP Allocation
13 (variable) **303V2**. Then continue by transitioning back to operating at the context as illustrated
14 Fig. 3A of G% data entropy and System OP Allocation (variable) **303V1**. The system OP is
15 dynamically varied from System OP Allocation **304G1** to System OP Allocation **304H2** and
16 then back to System OP Allocation **304G1**, in response to changes in, e.g., host allocation (e.g.
17 related to data entropy), system data write bandwidth, host data write bandwidth, and/or ratios
18 thereof.

19

20 **[0146]** Other embodiments are contemplated where allocation of OP resources between
21 host OP and system OP uses is based on data patterns as received from a host. For example, if a
22 host is performing a series of writes to sequential addresses, then an allocation between host and
23 system OP uses is computed using a reduced host data rate that is a fraction (less than unity) of a
24 measured host data rate, so that less host data OP is allocated than if the measured host data rate
25 were used for the allocation. For another example, if a host is performing a series of writes to
26 random addresses, then allocation between host and system OP uses is computed using an
27 increased system data rate that is a multiple of a measured system data rate, so that more system
28 data OP is allocated than if the measured system data rate were used for the allocation. The
29 multiple is greater than unity, but not necessarily an integer.

30

31 **[0147]** Other embodiments are contemplated where allocation of OP resources between
32 host OP and system OP uses is based on instantaneous usage of units of flash memory. For
33 example, if a burst of units are used to store system (or alternatively host) data, then an
34 allocation of system (or alternatively host) OP is temporarily boosted.

35

1 [0148] Some aspects of some embodiments represented by Fig. 3A and Fig. 3B (as well
2 as Fig. 2) are illustrated conceptually. In the aforementioned figures, the various allocations are
3 representative of respective total amounts of storage of a flash memory allocated, whether
4 contiguous or not. For example, System Allocation **201** is a plurality of non-contiguous units of
5 the Flash Memory. For another example, Host Allocation **302G1** and (combined) OP Allocation
6 **303G1** are each respective pluralities of non-contiguous units of the Flash Memory.

7

8 [0149] In some embodiments (such as illustrated in part by Figs. 2, 3A, and/or 3B),
9 additional flash memory (not illustrated) is available as a resource for increasing OP, or
10 alternatively failure of one or more portions of the flash memory results in decreasing OP. In
11 some embodiments, flash memory is held aside for uses not illustrated. For example, one or
12 more portions of one or more flash die of a flash memory are reserved for replacement of failed
13 portions of the flash memory.

14

15 [0150] In some embodiments (such as illustrated in part by Figs. 2, 3A, and/or 3B), all
16 or any portions of the host data and/or the system data are stored in a flash memory unaltered,
17 compressed, deduplicated, encrypted, subject to any form of reversible transformation, or any
18 combination thereof. In some embodiments, information stored in the flash memory includes, in
19 some embodiments, lower-level redundancy information (e.g. per-page ECC) and/or higher-
20 level redundancy information (e.g. RAID-like redundancy information), in addition to
21 information protected by the lower-level and/or higher-level redundancy information.

22

23

24 OVER-PROVISIONING (OP) RESOURCE UNIT (RE)ALLOCATION AND LIFE CYCLE

25

26 [0151] Fig. 4 illustrates Flow Diagram **400** of an embodiment of allocation and/or
27 (re)allocation of OP resources in a context of variable OP for managing NVMs, such as in one or
28 more contexts relating to any of Figs. 2, 3A, and 3B. Flow begins by computing if
29 circumstances correspond to performing a reallocation (Determine If (Re)Allocation Conditions
30 Exist **401** followed by (Re)Allocate? **402**). If the circumstances are not present, then flow loops
31 back to repeat the computing. If the circumstances are present, then a (re)allocation is requested,
32 queued, or immediately executed, according to various embodiments (Perform (Re)Allocation
33 **403**).

34

35 [0152] In various embodiments, a computation of whether allocation and/or
36 (re)allocation conditions exist (e.g. Determine If (Re)Allocation Conditions Exist **401** and/or

1 Perform (Re)Allocation **403**) are performed at any one or more of various particular points in
2 time. Examples of the various particular points in time are when one or more flash memory
3 units are allocated or marked as for a particular use (e.g., as being for system OP or host OP
4 use). Other examples are when the units exit, enter, or are picked for various operations such as
5 garbage collection, recycling, or erasing. Other examples are when the units enter a host unit
6 state, a system unit state, or a free queue, or when the units transition from one management
7 state to another. Other examples are when a number of the units available for a particular use
8 (e.g. for system use or host use) reach respective predetermined and/or programmatically
9 determined thresholds (e.g. high or low watermarks).

10

11 [0153] In various embodiments conceptually represented by any one or more of Figs. 2,
12 3A, 3B, and 4, a computation of whether allocation and/or (re)allocation conditions exist (e.g.
13 Determine If (Re)Allocation Conditions Exist **401**) is implemented as a unit is marked as being
14 for a particular usage (e.g. host or system).

15

16 [0154] In various embodiments conceptually represented by any one or more of Figs. 2,
17 3A, 3B, and 4, an allocation change, e.g. an allocation and/or a (re)allocation), is implemented as
18 a target when marking unused and/or empty units of the flash memory for a particular use, e.g.
19 marking a flash memory unit exiting garbage collection as being for host OP use versus system
20 OP use, or marking a flash memory unit at a time of entry into a free queue as being for host OP
21 use versus system OP use. For example, in some embodiments represented by Fig. 4, the
22 performing of reallocation is a multi-stage process. A target (re)allocation is determined and
23 stored, and then as flash memory units exit garbage collection, the exiting units are marked as
24 usable in accordance with the stored (re)allocation.

25

26 [0155] Fig. 5 illustrates Flow Diagram **500** of an embodiment of a life-cycle of a
27 managed unit of flash memory, such as an R-block. All or any portions of the flash memory are
28 managed as a collection of units, and the flow diagram is representative of how a flash unit
29 management agent manages state of each of the units. Conceptually Flow Diagram **500** is
30 representative of a state machine that is instantiated for each of the units that are managed. For
31 brevity in the following description, terminology such as “setting” a unit to a particular state or
32 “marking” a unit as being in a particular state is shorthand for setting management state
33 associated with the unit to the particular state (or to one or more values indicative of the
34 particular state).

35

1 [0156] Continuing with the description, at initialization (e.g. in response to power-on
2 reset, reception of a reset command, or any combination thereof) all of the units of the flash
3 memory are set to an initial state indicating the units are available for use (Free **501**). In
4 response to a request to prepare a unit for eventual allocation, the unit is prepared for allocation
5 (e.g. at least in part via garbage collection, recycling, and/or erasing) and is set to indicate
6 availability for immediate allocation (Free Queue **502**). In some embodiments, availability for
7 immediate allocation corresponds to the unit being in one of one or more free queues, as
8 conceptually indicated by an allocation transition marking (Queue Allocation **512**).
9

10 [0157] Once in the one of the free queues, a unit is marked according to usage as host
11 usage (Host Allocation **513**) or as system usage (System Allocation **514**). The unit is then used
12 in accordance with the marking for host usage (Host Unit **503**) or for system usage (System Unit
13 **504**). In response to choosing a particular unit for reclamation via recycling or garbage
14 collection, the particular unit is reclaimed and is set to indicate that the particular unit is
15 available for use (Free **501**). The setting to indicate availability occurs whether the unit was
16 being used for host usage (Recycle Host Unit **515**) before being chosen for reclamation, or
17 whether the unit was being used for system usage (Recycle System Unit **516**) before being
18 chosen for reclamation.
19

20 [0158] In some embodiments, choosing a unit for reclamation is conceptually a two-
21 stage process. First, a selection of a type of unit to reclaim, such as a unit being used to store
22 host data (corresponding, e.g., to Host Unit **503**) or a unit being used to store system data
23 (corresponding, e.g., to System Unit **504**) is made. The selection is based on (dynamically)
24 variable OP, such as described with respect to Figs. 2, 3A, or 3B. Second, within at least a
25 portion of the units of the selected category, a selection for the unit to reclaim is made.
26

27 [0159] In some alternate embodiments, there are no free queues, and a unit is marked as
28 transitioning immediately from being available for use (Free **501**) to being used according to
29 usage as a host unit (Host Unit **503**) or as a system unit (System Unit **504**) without passage
30 through an intermediate queuing state (Free Queue **502**). In some alternate embodiments, there
31 are two categories of queues according to units to be used as host units or system units, and a
32 decision about whether a unit is to be used as a host unit or as a system unit is made when the
33 unit is marked as transitioning immediately from being available for use (Free **501**) to being
34 available for immediate allocation according to a queue having the host categorization or the
35 system categorization.
36

1 [0160] In various embodiments and/or usage scenarios, various operations related to
2 any one or more of Figs. 2, 3A, 3B, and 4 are performed in response to and/or in coordination
3 with various transitions illustrated in Fig. 5 (or alternate embodiments described in relation to
4 Fig. 5). The various operations include the (dynamic) allocation(s), (dynamic) reallocation(s), as
5 well as computations and determinations relating to the allocations/reallocations.

6

7 [0161] For example, dynamically altering host allocation, such as illustrated between
8 Host Allocation **202A** and Host Allocation **202C** of Fig. 2, is performed when a unit is marked
9 as exiting the one of the free queues to be used as a host unit (Host Allocation **513**). For another
10 example, dynamically altering allocation of OP resources between host OP allocation and
11 system OP allocation, such as illustrated between System OP Allocation **204A** and System OP
12 Allocation **204B** of Fig. 2, is performed when a unit is marked as entering a free queue having a
13 system categorization. For yet another example, dynamically altering allocation of OP resources
14 between host OP allocation and system OP allocation, such as illustrated between System OP
15 Allocation **304G1** of Fig. 3A and System OP Allocation **304G2** of Fig. 3B, is performed as a
16 unit is marked as being for host use (Host Unit **503**) or system use (System Unit **504**) upon
17 exiting the one of the free queues.

18

19 [0162] For yet another example, computing whether or not to reallocate (e.g. Determine
20 If (Re)Allocation Conditions Exist **401** of Fig. 4), is performed in coordination with a request
21 that a unit be selected for reclamation and recycled (e.g. Recycle Host Unit **515** or Recycle
22 System Unit **516**). In some embodiments, a categorization of a unit selected for reclamation is
23 independent of a result of a (re)allocation (e.g. Perform (Re)Allocation **403**), such as a host unit
24 or a system unit being selected irrespective of a result of a reallocation. In other embodiments, a
25 categorization of a unit selected for reclamation is based at least in part on a result of a
26 (re)allocation (e.g. Perform (Re)Allocation **403**), such as a host unit being selected when a
27 reallocation result indicates that fewer host units are to be allocated to host OP usage.

28

29 [0163] In various embodiments, one or more elements of Figs. 2, 3A, 3B, and 4
30 correspond to or are related to one or more elements of Fig. 1A. For example, Flash Memory
31 **200** of Fig. 2 corresponds to NVM **199**. For another example, length along Physical Size **310** of
32 Fig. 3A and Fig. 3B corresponds to amount of storage of NVM **199**. For yet another example,
33 one or more of the allocation operations or allocation-related operations described with respect
34 to Figs. 2, 3A, 3B, and/or 4 are performed by, or under control of, one or more portions of one or
35 more of recycler **151** and CPU **171**. For yet another example, one or more of the state
36 transitions of Fig. 5 are performed by, or under control of, one or more portions of one or more

1 of map 141, recycler 151, and CPU 171. For yet another example, measurement of host data
2 rate(s), such as referred to with respect to Figs. 2, 3A, or 3B, is performed via all or any
3 portion(s) of host interface 111. For yet another example, measurement of host and/or system
4 data rate(s), such as referred to with respect to Figs. 2, 3A, or 3B, is performed via all or any
5 portion(s) of device interface logic 191 and/or scheduling 193.

6

7

8 EXAMPLE IMPLEMENTATION TECHNIQUES

9

10 [0164] In some embodiments, various combinations of all or portions of operations
11 performed by a system implementing variable OP for managing non-volatile storage, e.g. with
12 flash memories, a computing-host flash memory controller, and/or an SSD controller (such as
13 SSD controller 100 of Fig. 1A), and portions of a processor, microprocessor, system-on-a-chip,
14 application-specific-integrated-circuit, hardware accelerator, or other circuitry providing all or
15 portions of the aforementioned operations, are specified by a specification compatible with
16 processing by a computer system. The specification is in accordance with various descriptions,
17 such as hardware description languages, circuit descriptions, netlist descriptions, mask
18 descriptions, or layout descriptions. Example descriptions include: Verilog, VHDL, SPICE,
19 SPICE variants such as PSpice, IBIS, LEF, DEF, GDS-II, OASIS, or other descriptions. In
20 various embodiments, the processing includes any combination of interpretation, compilation,
21 simulation, and synthesis to produce, to verify, or to specify logic and/or circuitry suitable for
22 inclusion on one or more integrated circuits. Each integrated circuit, according to various
23 embodiments, is designable and/or manufacturable according to a variety of techniques. The
24 techniques include a programmable technique (such as a field or mask programmable gate array
25 integrated circuit), a semi-custom technique (such as a wholly or partially cell-based integrated
26 circuit), and a full-custom technique (such as an integrated circuit that is substantially
27 specialized), any combination thereof, or any other technique compatible with design and/or
28 manufacturing of integrated circuits.

29

30 [0165] In some embodiments, various combinations of all or portions of operations as
31 described by a computer readable medium having a set of instructions stored therein, are
32 performed by execution and/or interpretation of one or more program instructions, by
33 interpretation and/or compiling of one or more source and/or script language statements, or by
34 execution of binary instructions produced by compiling, translating, and/or interpreting
35 information expressed in programming and/or scripting language statements. The statements are
36 compatible with any standard programming or scripting language (such as C, C++, Fortran,

1 Pascal, Ada, Java, VBscript, and Shell). One or more of the program instructions, the language
2 statements, or the binary instructions, are optionally stored on one or more computer readable
3 storage medium elements. In various embodiments some, all, or various portions of the program
4 instructions are realized as one or more functions, routines, sub-routines, in-line routines,
5 procedures, macros, or portions thereof.

CONCLUSION

3 [0166] Certain choices have been made in the description merely for convenience in
4 preparing the text and drawings and unless there is an indication to the contrary the choices
5 should not be construed per se as conveying additional information regarding structure or
6 operation of the embodiments described. Examples of the choices include: the particular
7 organization or assignment of the designations used for the figure numbering and the particular
8 organization or assignment of the element identifiers (the callouts or numerical designators, e.g.)
9 used to identify and reference the features and elements of the embodiments.

11 [0167] The words “includes” or “including” are specifically intended to be construed as
12 abstractions describing logical sets of open-ended scope and are not meant to convey physical
13 containment unless explicitly followed by the word “within.”

15 [0168] Although the foregoing embodiments have been described in some detail for
16 purposes of clarity of description and understanding, the invention is not limited to the details
17 provided. There are many embodiments of the invention. The disclosed embodiments are
18 exemplary and not restrictive.

20 [0169] It will be understood that many variations in construction, arrangement, and use
21 are possible consistent with the description, and are within the scope of the claims of the issued
22 patent. For example, interconnect and function-unit bit-widths, clock speeds, and the type of
23 technology used are variable according to various embodiments in each component block. The
24 names given to interconnect and logic are merely exemplary, and should not be construed as
25 limiting the concepts described. The order and arrangement of flowchart and flow diagram
26 process, action, and function elements are variable according to various embodiments. Also,
27 unless specifically stated to the contrary, value ranges specified, maximum and minimum values
28 used, or other particular specifications (such as flash memory technology types; and the number
29 of entries or stages in registers and buffers), are merely those of the described embodiments, are
30 expected to track improvements and changes in implementation technology, and should not be
31 construed as limitations.

33 [0170] Functionally equivalent techniques known in the art are employable instead of
34 those described to implement various components, sub-systems, operations, functions, routines,
35 sub-routines, in-line routines, procedures, macros, or portions thereof. It is also understood that
36 many functional aspects of embodiments are realizable selectively in either hardware (e.g.,

1 generally dedicated circuitry) or software (e.g., via some manner of programmed controller or
2 processor), as a function of embodiment dependent design constraints and technology trends of
3 faster processing (facilitating migration of functions previously in hardware into software) and
4 higher integration density (facilitating migration of functions previously in software into
5 hardware). Specific variations in various embodiments include, but are not limited to:
6 differences in partitioning; different form factors and configurations; use of different operating
7 systems and other system software; use of different interface standards, network protocols, or
8 communication links; and other variations to be expected when implementing the concepts
9 described herein in accordance with the unique engineering and business constraints of a
10 particular application.

11

12 [0171] The embodiments have been described with detail and environmental context
13 well beyond that required for a minimal implementation of many aspects of the embodiments
14 described. Those of ordinary skill in the art will recognize that some embodiments omit
15 disclosed components or features without altering the basic cooperation among the remaining
16 elements. It is thus understood that much of the details disclosed are not required to implement
17 various aspects of the embodiments described. To the extent that the remaining elements are
18 distinguishable from the prior art, components and features that are omitted are not limiting on
19 the concepts described herein.

20

21 [0172] All such variations in design are insubstantial changes over the teachings
22 conveyed by the described embodiments. It is also understood that the embodiments described
23 herein have broad applicability to other computing and networking applications, and are not
24 limited to the particular application or industry of the described embodiments. The invention is
25 thus to be construed as including all possible modifications and variations encompassed within
26 the scope of the claims of the issued patent.

WHAT IS CLAIMED IS:

- 1 1. A system comprising:
 - 2 a means for operating all or any portions of one or more flash memories as respective
 - 3 allocations dedicated respectively to host data, system data, system over-
 - 4 provisioning (OP), and host OP;
 - 5 a means for dynamically determining any one or more of the respective allocations in
 - 6 response to one or more events;
 - 7 wherein the means for operating and the means for dynamically determining are
 - 8 comprised in a controller of a Solid-State Disk (SSD) and the flash memories
 - 9 are comprised in the SSD; and
 - 10 wherein the events comprise a change in a used amount of the host data allocation due at
 - 11 least in part to a change in compression of host data stored in the host data
 - 12 allocation.
- 1 2. The system of claim 1, wherein the means for dynamically determining increases one or more
- 2 of the system data allocation and the system OP allocation in response to a reduction in the used
- 3 amount of the host data allocation due to improved compression of the used amount of the host
- 4 data allocation.
- 1 3. The system of claim 1, wherein the means for dynamically determining increases one or more
- 2 of the system OP allocation and the host OP allocation in response to a reduction in the used
- 3 amount of the host data allocation due to improved compression of the used amount of the host
- 4 data allocation.
- 1 4. The system of claim 3, wherein the means for dynamically determining allocates between the
- 2 system OP allocation and the host OP allocation based at least in part on respective bandwidths
- 3 of writes to the system data allocation and to the host data allocation.
- 1 5. The system of claim 1, wherein the change in compression is determined at least in part by
- 2 determining a change in amount of the host data allocation that is unused.
- 1 6. The system of claim 1, wherein the change in the used amount of the host data allocation is
- 2 due at least in part to a change in deduplication of the stored host data.

1 7. The system of claim 1, wherein the change in the used amount of the host data allocation is
2 due at least in part to a host command that explicitly de-allocates all or any portions of the stored
3 host data.

1 8. The system of claim 1, wherein changes to any one or more of the respective allocations is
2 effective any one or more of immediately, after garbage collecting, after recycling, and after
3 erasing.

1 9. The system of claim 1, further comprising a means for interfacing with a host, the means for
2 interfacing comprised in the controller and enabled to receive data from the host.

1 10. A method comprising:
2 operating all or any portions of one or more flash memories as respective allocations
3 dedicated respectively to host data, system data, system over-provisioning (OP),
4 and host OP;
5 dynamically determining any one or more of the respective allocations in response to
6 one or more events;
7 wherein the operating and the dynamically determining are performed by a controller
8 comprised in a Solid-State Disk (SSD); and
9 wherein the events comprise a change in a used amount of the host data allocation due at
10 least in part to a change in compression of host data stored in the host data
11 allocation.

1 11. The method of claim 10, wherein the dynamically determining increases one or more of the
2 system data allocation and the system OP allocation in response to a reduction in the used
3 amount of the host data allocation due to improved compression of the used amount of the host
4 data allocation.

1 12. The method of claim 10, wherein the dynamically determining increases one or more of the
2 system OP allocation and the host OP allocation in response to a reduction in the used amount of
3 the host data allocation due to improved compression of the used amount of the host data
4 allocation, and allocates between the system OP allocation and the host OP allocation based at
5 least in part on respective bandwidths of writes to the system data allocation and to the host data
6 allocation.

1 13. The method of claim 10, wherein the change in compression is determined at least in part by
2 determining a change in amount of the host data allocation that is unused.

1 14. The method of claim 10, wherein the change in the used amount of the host data allocation
2 is due at least in part to one or more of a change in deduplication of the stored host data and a
3 host command that explicitly de-allocates all or any portions of the stored host data.

1 15. The method of claim 10, wherein changes to any one or more of the respective allocations is
2 effective any one or more of immediately, after garbage collecting, after recycling, and after
3 erasing.

1 16. A tangible computer readable medium having a set of instructions stored therein that when
2 executed by a processing element cause the processing element to perform and/or control
3 operations comprising:

4 managing all or any portions of one or more flash memories as respective allocations
5 dedicated respectively to host data, system data, system over-provisioning (OP),
6 and host OP;
7 dynamically determining any one or more of the respective allocations in response to
8 one or more events;
9 wherein the tangible computer readable medium and the processing element are
10 comprised in a Solid-State Disk (SSD); and
11 wherein the events comprise a change in a used amount of the host data allocation due at
12 least in part to a change in compression of host data stored in the host data
13 allocation.

1 17. The tangible computer readable medium of claim 16, wherein the dynamically determining
2 increases one or more of the system OP allocation and the host OP allocation in response to a
3 reduction in the used amount of the host data allocation due to improved compression of the
4 used amount of the host data allocation, and allocates between the system OP allocation and the
5 host OP allocation based at least in part on respective bandwidths of writes to the system data
6 allocation and to the host data allocation.

1 18. The tangible computer readable medium of claim 16, wherein the change in compression is
2 determined at least in part by determining a change in amount of the host data allocation that is
3 unused.

1 19. The tangible computer readable medium of claim 16, wherein the change in the used
2 amount of the host data allocation is due at least in part to one or more of a change in
3 deduplication of the stored host data and a host command that explicitly de-allocates all or any
4 portions of the stored host data.

1 20. The tangible computer readable medium of claim 16, wherein changes to any one or more of
2 the respective allocations is effective any one or more of immediately, after garbage collecting,
3 after recycling, and after erasing.

1/7




Fig. 1A

2/7




Fig. 1B

3/7




Fig. 2

4/7




Fig. 3A

5/7




Fig. 3B

6/7

400




Fig. 4

7/7




Fig. 5