

Exhibit A



US008024395B1

(12) **United States Patent**
Odom

(10) **Patent No.:** **US 8,024,395 B1**
(45) **Date of Patent:** **Sep. 20, 2011**

(54) **DISTRIBUTED PROCESSING MULTIPLE TIER TASK ALLOCATION**

(76) Inventor: **Gary Odom**, Portland, OR (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 508 days.

(21) Appl. No.: **10/228,588**

(22) Filed: **Aug. 26, 2002**

Related U.S. Application Data

(60) Provisional application No. 60/317,108, filed on Sep. 4, 2001.

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/201**; 709/223; 709/226

(58) **Field of Classification Search** 705/37; 709/226, 201, 208, 223; 718/105

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,662,401	A *	5/1972	Collins et al.	718/103
5,025,369	A *	6/1991	Schwartz	718/102
6,112,225	A *	8/2000	Kraft et al.	709/202
6,167,427	A *	12/2000	Rabinovich et al.	709/201
6,192,388	B1 *	2/2001	Cajolet	718/100
6,263,358	B1 *	7/2001	Lee et al.	718/100
6,370,560	B1 *	4/2002	Robertazzi et al.	718/105
6,463,457	B1 *	10/2002	Armentrout et al.	709/201

6,775,831	B1 *	8/2004	Carrasco et al.	707/200
6,782,422	B1 *	8/2004	Bahl et al.	709/224
6,826,753	B1 *	11/2004	Dageville et al.	718/102
7,188,113	B1 *	3/2007	Thusoo	1/1
7,383,426	B2 *	6/2008	Chung et al.	712/220
7,668,800	B2 *	2/2010	Motoyama et al.	707/999.001
7,849,178	B2 *	12/2010	Shen et al.	709/223
2004/0045002	A1 *	3/2004	Berger et al.	718/102
2004/0264503	A1 *	12/2004	Draves, Jr.	370/469
2008/0216859	A1 *	9/2008	Chan	132/224
2009/0204470	A1 *	8/2009	Weyl et al.	705/9

OTHER PUBLICATIONS

Kao, "Subtask deadline assignment for complex distributed soft real-time tasks," Proceedings of the 14th International Conference on Distributed Computing Systems, 1994; Jun. 21-24, 1994; pp. 172-181, USA.

Lee, "Some simple task assignment problems for distributed networked agents," Fourth International Conference on Knowledge-Based Intelligent Engineering Systems and Allied Technologies, 2000. Proceedings. vol. 1, Aug. 30-Sep. 1, 2000 pp. 305-308, USA.

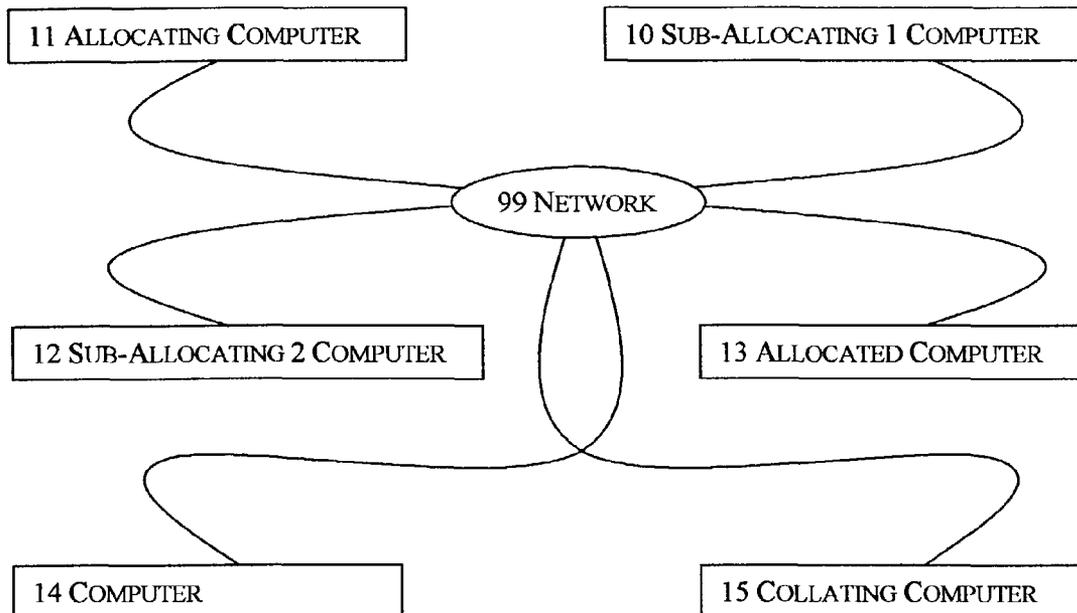
* cited by examiner

Primary Examiner — Faruk Hamza

(57) **ABSTRACT**

Described is a system and methods for multiple tier distribution of task portions for distributed processing. Essentially, a task is divided into portions by a first computer and a task portion transferred to second participatory computer on the network, whereupon an allocated task portion is again portioned by the second computer into subtask portions, and a subtask portion transferred by the second computer to a third participatory computer on the network, whereby distributed processing transpires, and results collated as required.

20 Claims, 6 Drawing Sheets



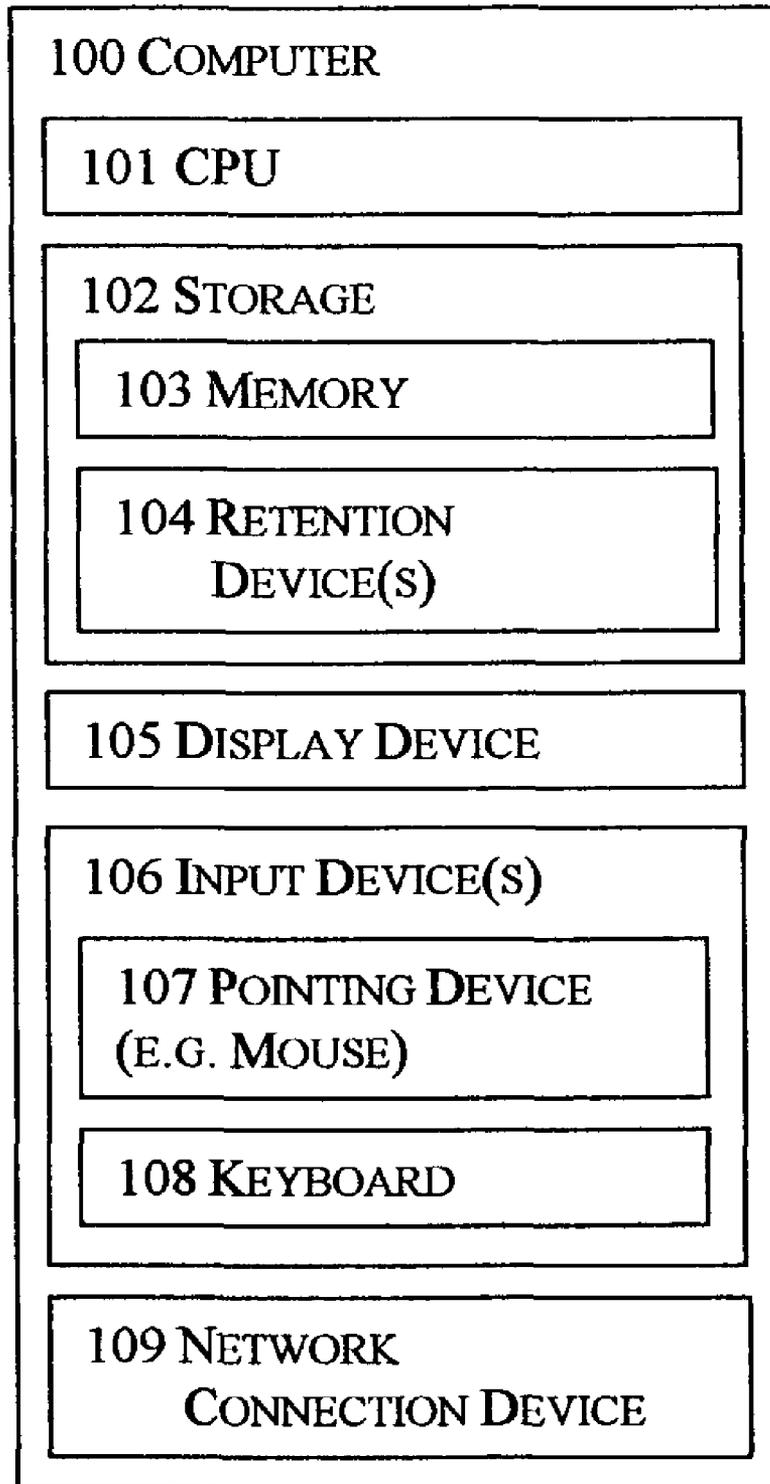


FIGURE 1

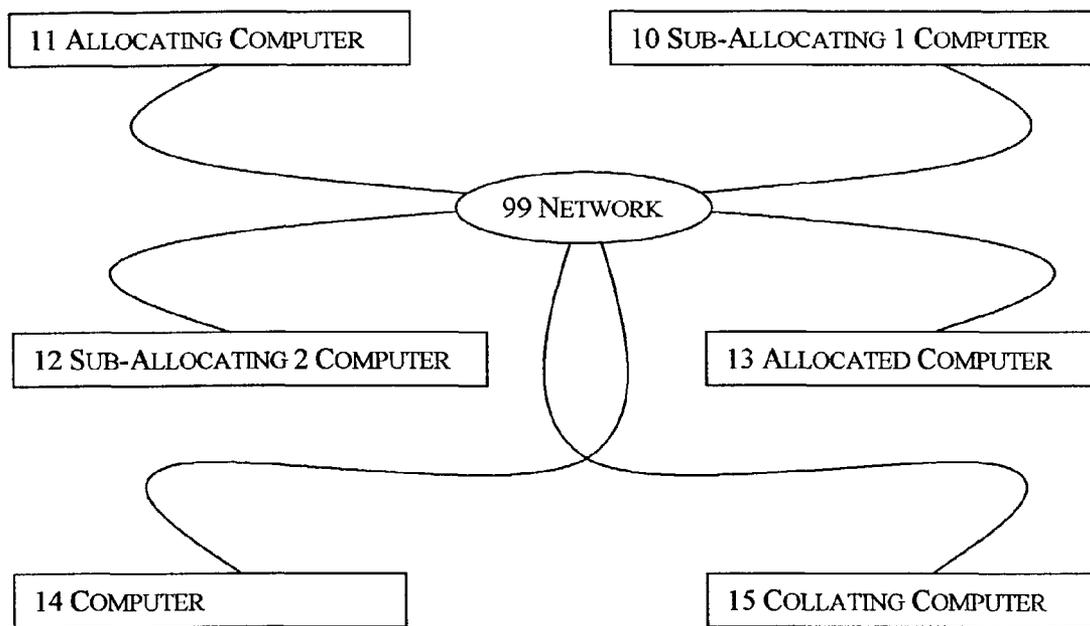


FIGURE 2

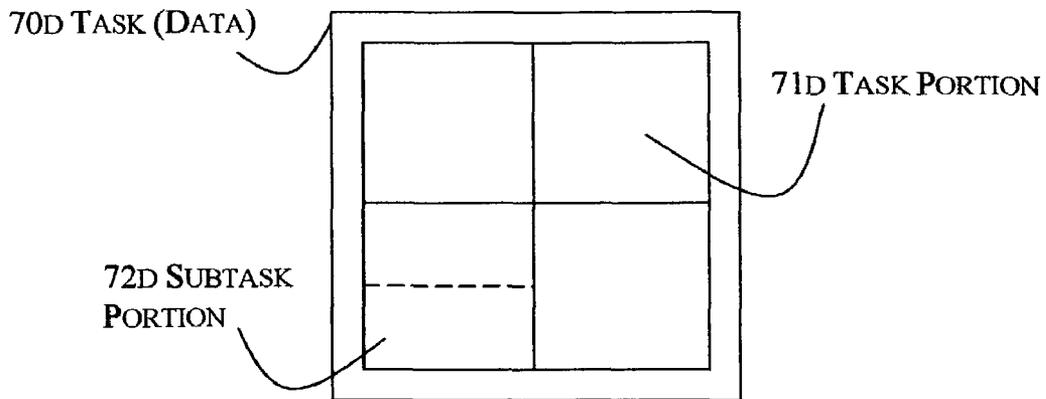


FIGURE 3A

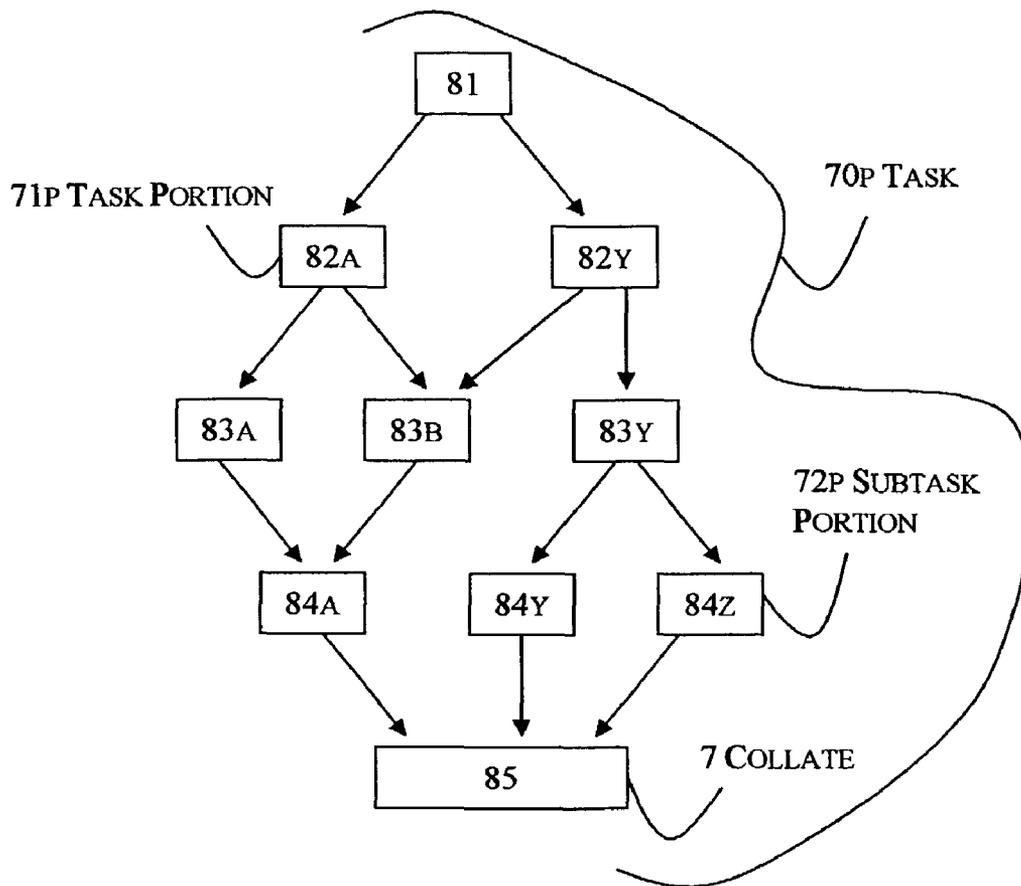


FIGURE 3B

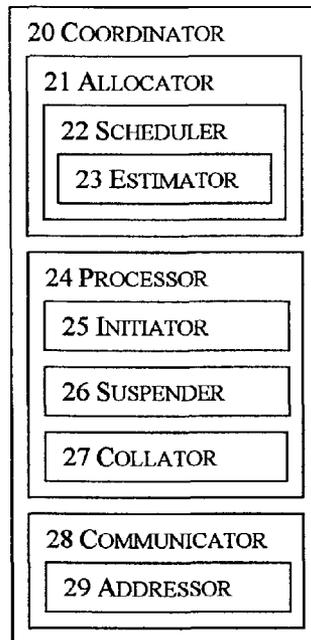


FIGURE 4

60 (SUB)TASK PORTION	61 MESSAGE TYPE	62 DATA	63 CODE	64 STATUS / DIRECTIVE	65 RESULTS
IDENTIFIABLE BY DIVISION	ACTION-SPECIFIC (E.G. ALLOCATE, CANCEL, RESULT, COLLATE)	(SUB)TASK DATA OR DATA REFERENCE	SOFTWARE OR SW REFERENCE	DEPENDS ON MESSAGE TYPE	DEPENDS ON MESSAGE TYPE

FIGURE 5

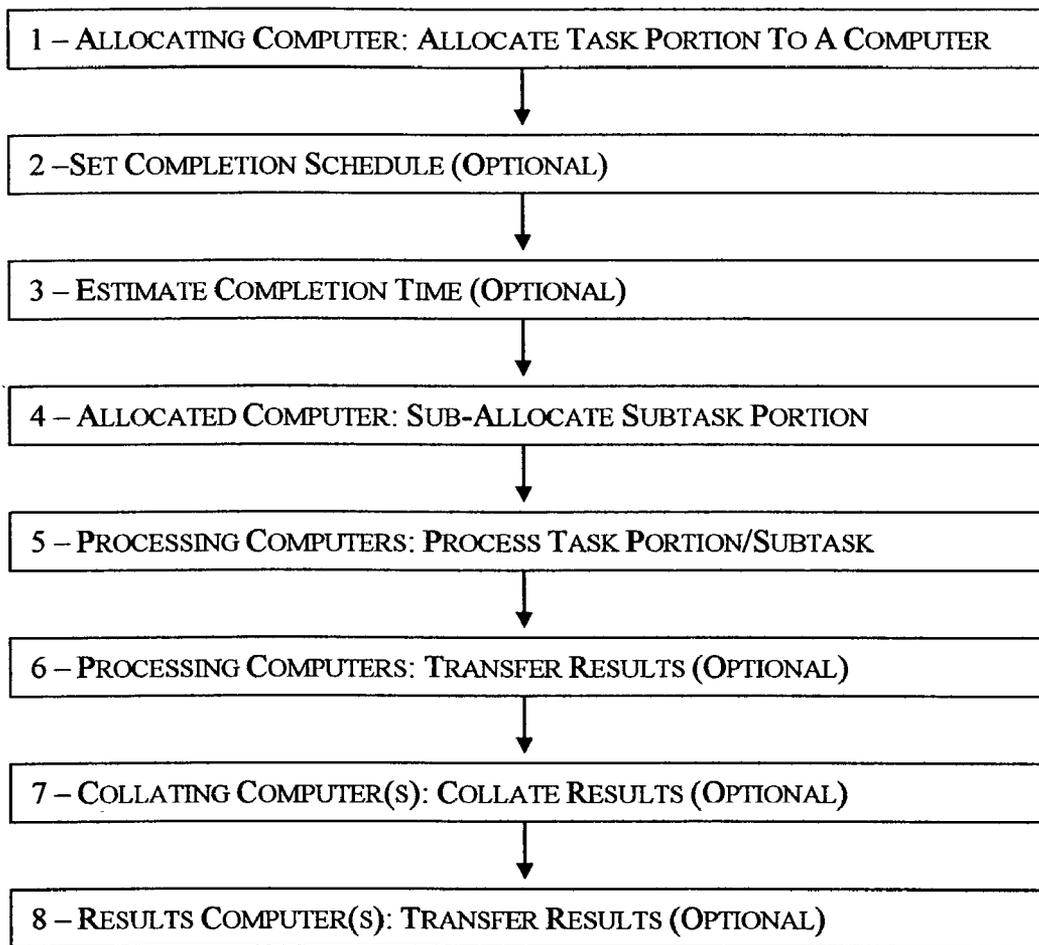


FIGURE 6

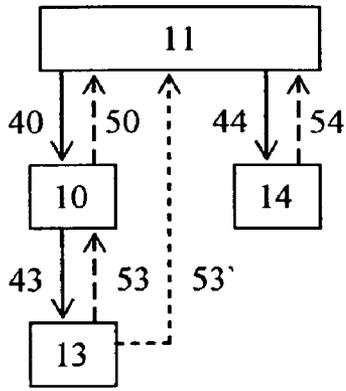


FIGURE 7A

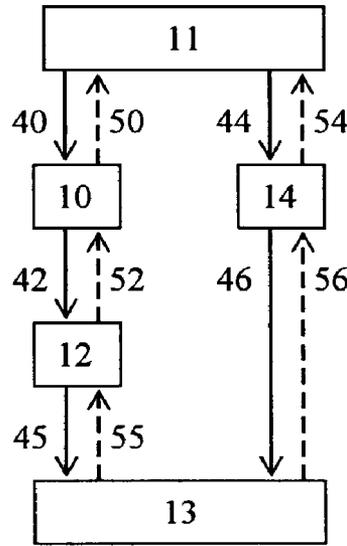


FIGURE 7B

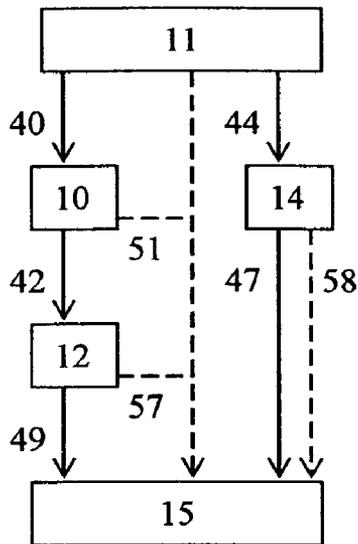


FIGURE 7C

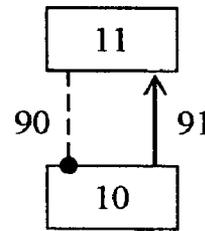
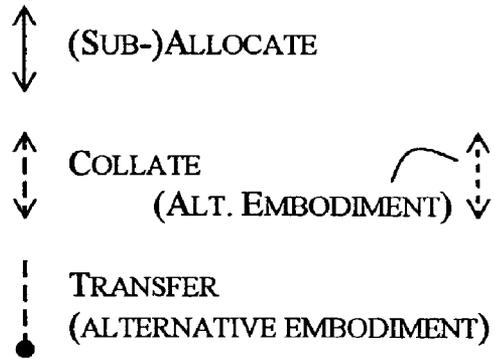


FIGURE 7D

1

**DISTRIBUTED PROCESSING MULTIPLE
TIER TASK ALLOCATION**CROSS-REFERENCE TO RELATED
APPLICATIONS

This application **claims priority benefit under 35 U.S.C. §119(e) of U.S. Provisional Application No. 60/317,108, filed Sep. 4, 2001.**

STATEMENT REGARDING FEDERALLY
SPONSORED RESEARCH OR DEVELOPMENT

Not Applicable

THE NAMES OF THE PARTIES TO A JOINT
RESEARCH AGREEMENT

Not Applicable

INCORPORATION-BY-REFERENCE OF
MATERIAL SUBMITTED ON A COMPACT DISC

Not Applicable

BACKGROUND OF THE INVENTION

1. Field of the Invention

The relevant technical field is computer software, specifically distributed processing in a networked environment.

2. Description of the Related Art

Including Information Disclosed Under 37 CFR 1.97 and 1.98

In what is not ironically called a “network effect”, the advantage of distributed processing is positively correlated to availability of powerful computers in a networked environment. This trend is especially encouraged by always-on broadband connection to the ultimate wide-area network: the Internet.

U.S. Pat. No. 6,192,388 details “detecting available computers to participate in computationally complex distributed processing problems”, and switching an allocated task portion to a different computer if the one first assigned the task portion becomes occupied. U.S. Pat. No. 6,192,388 also describes some of the resource factors involved in determining whether to allocate a task portion to a computer.

With some content overlap to the earlier-filed U.S. Pat. No. 6,192,388, U.S. Pat. No. 6,112,225 describes a “task distribution processing system and the method for subscribing computers to perform computing tasks during idle time”, and goes into detail as to various ways of specifying “idle time”. Both U.S. Pat. Nos. 6,192,388 and 6,112,225, incorporated herein by reference, use the same computer for allocating, monitoring and re-allocating task portions.

U.S. Pat. No. 6,263,358 describes sophisticated regimes of scheduling of distributed processing tasks using software agents. In the face of schedule slippage, such a system relies upon coordination among multiple agents to work effectively.

BRIEF SUMMARY OF THE INVENTION

Multiple tier task allocation maximizes flexibility and productivity of distributed processing participatory computers.

A computer which has been allocated a distributed processing task portion may itself determine to reallocate a portion of its subtask, for example, in order to meet a schedule, or if its performance profile deteriorates below expectation. The

2

described technology localizes further (sub)task portion allocation control to computers having been assigned task portions.

Further task processing division to other computers on the network may be extended to initial task portioning, scheduling, and results collation.

Admittedly, only those tasks capable of being subdivided in some manner may benefit from the described technology.

BRIEF DESCRIPTION OF THE SEVERAL
VIEWS OF THE DRAWINGS

FIG. 1 is a block diagram of a suitable computer.

FIG. 2 depicts an example computer network.

FIG. 3 depicts example tasks.

FIG. 4 depicts relevant distributed processing application components.

FIG. 5 depicts an abstract of a distributed processing message.

FIG. 6 depicts distributed processing steps.

FIG. 7 depicts examples of processing distribution and results collation.

DETAILED DESCRIPTION OF THE INVENTION

25

FIG. 1 is a block diagram of a computer **100** which comprises at least a CPU **101**; storage **102**, which comprises memory **103** and optionally one or more devices with retention medium(s) **104** such as hard disks, diskettes, compact disks (e.g. CD-ROM), or tape; a device **109** for connection to a network **99**; an optional display device **105**; and optionally one or more input devices **106**, examples of which include but are not exclusive to, a keyboard **108**, and/or one or more pointing devices **107**, such as a mouse. Such a computer **100** is suitable for the described technology.

30

FIG. 2 is a block diagram of distributed processing participatory computers **100** connected to each other through a network **99**. Computers **100** are participatory based upon having installed required software and, optionally, meeting specified conditions for participation. Example conditions include sufficient processing power, storage, network bandwidth or reliability, or adequate security precautions, such as a particular installed operating system.

40

Computer **11** in FIG. 2 is depicted in the role of an allocating computer, signifying initial allocation of task portions. Likewise, other computers in FIG. 2 are signified by their roles. FIGS. 2, 6, and 7 are used for example explanation of the technology. **The roles of computers are envisioned as transitory: for example, a computer initiating distributed processing and allocating task portions for its task may next have a task or sub-task portion allocated to it by another computer in a succeeding task.**

45

50

A network **99** may be any means by which computers are connected for software program or data transfer. The described technology relies upon network connectivity, including inter-application messaging and data or software transfer capabilities that are well known.

55

Participatory computers have software installed enabling the desired distributed processing. The software may be installed by download through network **99** connection, or via a more traditional local retention medium, such as CD-ROM or floppy disk.

60

The desired distributing processing may take various forms. FIG. 3 illustrates examples.

65

One example is a divisible and distributable chunk of data requiring a single processing, as depicted in FIG. 3a, split into portions so that the various participatory computers can pro-

cess the data portions. The task data **70d** is shown portioned into equal quarter task portions **71d**. A task portion has been further split into subtask portions **72d**.

An example at the other end of the spectrum, depicted in FIG. **3b**, is a series of processing steps which to some extent may overlap, whereby each of the participatory computers performs some portion of the task **70**. Task **70p** processing can be portioned into task portions **71p** (**82a-84a** and **82y-84y/z**). Further, a subtask portion **72p** could be allocated at specific processing steps (**83a/b** or **84y/z**). Note that synchronization may be an issue, such as in FIG. **3b** where processing step **83b** requires the output of preceding steps **82a** and **82y** to proceed. There may also be a results collation **85** step. Between the extreme examples lies divisible and distributable data capable of being processed in an overlap (not exclusively serial) manner.

One possible employment scenario for the described technology is a set of participatory computers running one or more applications which intermittently require intermittent excessive (to a single computer) processing. Distributed processing may be used as a remedy for those times when a singular computer may otherwise bog down or be insufficient. In this scenario, any computer with excessive processing needs may initiate shared task processing, either by direct allocation of task portions, or by directing another computer to perform task portion allocation and attendant processing.

Note that the term "allocate" and its conjugations may refer to initial allocation or subsequent sub-allocation—after all, the allocation process is self-similar. In the preferred embodiment, allocation (and sub-allocation) necessarily implies portioning of a (sub)task prior to transferring a portion to another computer. In an alternative embodiment, depicted in FIG. **7d**, a task or (sub)task portion may be (sub-)allocated by transfer **90** of the (sub)task to another computer **10** prior to any portioning by the initially transferring computer **11**, with a request or directive that a portion be (sub-)allocated **91** to the computer **11** initiating the transfer, thus putting the overhead of (sub-)allocation on the recipient **10** rather than the initially transferring computer **11**.

FIG. **4** depicts an exemplary embodiment of relevant components of a distributed processing program, some of which are optional, depending upon embodiment; other components, such as user interface, event handling, and the actual processing modules, likely exist. Components may have different configurations in different embodiments.

While an application program is used as the preferred embodiment, an alternative preferred embodiment may incorporate all or portions of the described distributed processing functionality in an operating system.

An overall coordinator **20** may be employed to ensure proper interaction between the relevant distributed processing modules. In one embodiment, certain modules may be missing from an application on a particular computer, in which case the coordinator **20** would know the (limited) capabilities of the application, and compensate accordingly. Operationally, that compensation may take the form of knowing, by an addressor **29** with a database tracking such capabilities, of suitable computers with adequate capabilities to take on jobs which a coordinator **20** needs to off-load.

For example, a computer with limited storage or processing power may not have a scheduler **22** or collator **27**, whereby a coordinator **20** off-loads those jobs to an appropriate computer. A powerful computer with relative poor network capacity (speed or reliability) may be shunned from communication-intensive jobs, such as collation **7**. In this

embodiment, the distributed processing application may be heterogeneous, comprising relative capabilities according to computer capacity.

Messages are passed as required, including, for example, the following types of messages **61**: (sub)task portion allocation; data **62** or code **63** transfer; cancellation; scheduling; directives or estimation initiation or results; processing: directives (such as initiation, suspension, or collation) and results **65**. FIG. **5** depicts an abstract of a distributed processing message; intended for conceptual understanding and suggestion, not specific implementation (as this message format is not particularly efficient). Not all fields shown would necessarily be used for each message type **61**, and other fields may be required depending upon message type **61** or embodiment.

(Sub)task portions may be identifiable by its division, such as, for example: $2/5-1/4-2/3$, where each set of number indicates a (sub)task division. $2/5$, for example, would be part 2 of 5 portions. The point is to allow portioning by an allocator **21** and recombination of results by a collator **27**. A table or database may be kept and transferred as necessary that identifies actual and/or possible (sub)task portions.

Data **62** or executable software code **63** or references to them may be transferred via messaging. Status/directive **64** and result **65** depend on message type **61**.

Keeping track of processing times of allocated (sub)tasks (including CPU overhead and other performance factors) by computer is recommended as a way to calibrate future allocations.

FIG. **6** outlines the steps for the described multiple tier distributed processing. FIG. **7** illustrates examples of the distribution process.

An allocating computer **11** allocates a portion of a task to another computer **10** in step **1**. As depicted in FIG. **7a**, an allocating computer **11** may allocate task portions to multiple computers (**11** and **14**). An allocator **21** may be employed for task and subtask portioning and transfer, and for tracking such (sub-)allocations and portions.

Optionally, an allocating **11** (or sub-allocating **10**) or allocated **13** computer may set a completion schedule (step **2**) for the time by which results should be available. Depending upon the nature of the task, a schedule may be a single completion time for an allocated portion, or for intermediate computations as well. Ostensibly, a schedule is the raison d'être for multiple tier subtask sub-allocation, but subtask sub-allocation may be driven by anticipation of available resources which later fail to appear forthcoming. For example, an allocated computer **13** may become busier than historical usage would indicate, making (sub)task portion offloading prudent.

If scheduling is a factor, an estimated completion time calculation (step **3**) is advised. The availability and speed of resources, such as processor(s) **101** and storage **102**, may naturally figure into such calculation. Estimation calculations may be done by any participatory computer with sufficient information.

As depicted in FIG. **4**, an allocator **21** may employ a scheduler **22**, which may employ an estimator **23**, to perform processing steps **3** and **2** respectively.

The overhead of distribution may be considered by an estimator **23** or scheduler **22** as a factor in (sub-)allocation. Distribution overhead includes the time and resources to portion and distribute subtask portions, and to collect and collate results. Depending on the network, communication lags may also be a factor. Excessive (sub)task portion (sub-)allocation granularity is conceivable and should be accounted for. A suggested rule is that estimate of (sub-) allocation should be

5

a fraction of estimated processing time if processing time is the bottleneck; storage 102 capacity or other such bottlenecks necessitate similar consideration.

An estimate of processing capability may be ascertained for a computer targeted for processing prior to (sub-)allocation, so as to portion (sub)tasks accordingly.

For whatever reason, in step 4, a computer 10 with an allocated task portion 71 decides to sub-allocate a portion 72 of its allotted subtask to another computer 13, as depicted in FIG. 7a.

Participatory computers with (sub-)allocated (sub)task portions perform required processing per step 5. The generic processor 24 signifies the performer of step 5. An initiator 25 may synchronize with other processors 24 if necessary. A computer may be watchful (a possible coordinator 20 job) and sub-allocate after beginning processing, upon realizing sub-allocation as a prudent measure because of some unanticipated constraint, such as, for example, high CPU utilization (processing overhead) or suddenly limited storage. A suspender 26 may suspend processing, saving state as necessary for later resumption.

Depending upon embodiment, processing may occur only under specified conditions, for example, only when a computer is past a threshold state deemed idle. Other conditions, such as available storage 102, or network 99 connection speed or reliability, may also be pertinent allocation or processing criteria. If processing is conditional, temporary results may be stashed (locally or elsewhere on the network) for later resumption. A processor 24 initiator 25 and suspender 26 may, for example, respectively detect and act upon onset and termination of specified threshold conditions.

Step 6 specifies transferring results. This step may not be necessary, depending upon the task 70. Likewise, in step 7, results are optionally collated by one or more participatory computers, with results monitoring as required. Results monitoring and collation may itself become a distributed task. Collators 27 on multiple computers may collaborate to piece together and conclude the task.

With the notable exception of 53', FIG. 7a depicts results returned to the computer which allocated (or sub-allocated) the task (subtask) portion (50, 53, 54) for collation. But, as shown by example, results may be sent 53' to the allocating computer 11 instead of or in addition to that computer 10 that (sub-)allocated a (sub)task portion.

FIG. 7c depicts results being transmitted (likely for collation) to a different computer 15 than the allocating computer 11. This strategy may make sense, for example, when a series of tasks are allocated in succession: a division of duty between an allocating computer 11 and a results-collating computer 15. Final results may be sent to the allocating computer 11 or other computers by the collating computer 15 as necessary.

FIG. 7b depicts a situation where an allocated computer 13 is processing multiple subtask portions allocated by different computers (12, 14). This is doable given identifiable portions as suggested.

Task or subtask portions may be redundantly assigned as a precaution. Redundant (sub)allocation may be sensible given scheduling constraints.

Security may be an issue. Data, results, messages, or other content may be encrypted as required.

The invention claimed is:

1. A computer-implemented method for distributed processing comprising:

dividing a task into a plurality of task portions, wherein said task comprises at least one of divisible data or divisible executable instruction sets;

6

an allocating computer transferring at least one said task portion to a sub-allocating computer;
said sub-allocating computer receiving said task portion;
said sub-allocating computer dividing said task portion into a plurality of subtask portions;
said sub-allocating computer transferring at least one said subtask portion to an allocated computer,
said allocated computer receiving said subtask portion;
said allocated computer processing said subtask portion, whereby producing at least one result;
said allocated computer transferring said result to a pre-designated results computer;
said results computer receiving and storing said result; and
such that all foregoing transferring occurs by network connection.

2. The method according to claim 1, wherein said sub-allocating computer conditionally determines allocating said task portion to said allocated computer.

3. The method according to claim 1, further comprising: said sub-allocating computer redundantly allocating said subtask portion.

4. The method according to claim 1, wherein said sub-allocating computer receives indicia of predetermined subtask portions of said task portion separate from receiving said task portion.

5. The method according to claim 1, further comprising: said sub-allocating computer dividing said task portion based upon, at least in part, a schedule and estimation related to said schedule.

6. The method according to claim 5, wherein said schedule is not received from said allocating computer.

7. A computer-implemented method for distributed processing comprising:

dividing a task into a plurality of task portions, wherein at least one first task portion comprises further divisible portions,

hereinafter referred to as subtask portions;

an allocating computer allocating said first task portion to a sub-allocating computer via network connectivity;
said sub-allocating computer receiving said first task portion;

said sub-allocating computer dividing said first task portion into a plurality of subtask portions;

said sub-allocating computer allocating at least one said subtask portion to an allocated computer via network connectivity;

said allocated computer receiving said subtask portion;
said allocated computer processing said subtask portion, whereby producing at least one result related to said subtask portion and storing said result;

said allocated computer transferring said result to a results computer;

said results computer receiving a plurality of results related to said first task; and

said results computer collating said results.

8. The method according to claim 7, further comprising: said sub-allocating computer communicating with said allocated computer regarding subtask allocation prior to allocating said subtask portion to said allocated computer.

9. The method according to claim 7, further comprising: said results computer receiving redundant results portions.

10. The method according to claim 7, further comprising: said sub-allocating computer determining said subtask portion allocation by relying partly upon a schedule.

7

11. The method according to claim 7, wherein said allocating computer and said results computer comprise the same computer.

12. A computer-implemented method for distributed processing comprising:

dividing a task into a plurality of task portions;

an allocating computer allocating at least one said task portion to a sub-allocating computer;

said sub-allocating computer receiving said task portion;

said sub-allocating computer allocating a subtask portion to an allocated computer,

wherein said subtask portion comprises a portion of a task portion;

said allocated computer receiving said subtask portion;

a subtask processing computer processing said subtask portion,

thereby creating and storing at least one result;

said subtask processing computer transferring said result to a results computer; and

said results computer receiving and storing results related to said task from a plurality of computers.

13. The method according to claim 12, wherein said sub-allocating computer partially processes said at least a portion of said task portion prior to allocating said subtask portion to said allocated computer.

8

14. The method according to claim 12, further comprising: said sub-allocating computer selecting said allocated computer based, at least in part, upon network communication with at least one other computer.

15. The method according to claim 12, further comprising: associating a schedule with said subtask portion.

16. The method according to claim 12, wherein said allocated computer and said subtask processing computer comprise the same computer.

17. The method according to claim 12, wherein said allocated computer and said subtask processing computer comprise different computers.

18. The method according to claim 12, wherein said allocated computer processes a plurality of subtask portions received from a plurality of computers.

19. The method according to claim 12, further comprising: conditionally determining at least one of a sub-allocating computer and an allocated computer based, at least in part, upon data received via network communication.

20. The method according to claim 12, further comprising: said allocated computer conditionally allocating said subtask portion to said subtask processing computer.

* * * * *