

CLAIMS

What is claimed is:

1. A computing system (104) including:

a distributed computing cluster (112) including:

a plurality of computing nodes (105) interconnected by an interconnect network (107) over which the computing nodes of the plurality of computing nodes communicate with each other by passing messages, each of at least some of the computing nodes being configured with a first parameter governing transmissions of messages by the computing node over the interconnect network,

wherein the at least some computing nodes are configured to accumulate messages for transmission as a group of messages according to the first parameter, and

wherein the computing system is configured to limit injections of computing requests into the distributed computing cluster according to a second parameter; and the computing system further includes

a controller (110) configured to receive at least one predetermined service level requirement and to control a value of the second parameter and a value of the first parameter to control a computational throughput of the distributed computing cluster while complying with the at least one service level requirement.

2. The computing system of claim 1 wherein the distributed computing cluster includes a rate limiter (117) configured to limit a rate of injection of the computing requests into the distributed computing system according to the second parameter.

3. The computing system of claim any one of claims 1 to 2 wherein each service level requirement of the one or more service level requirements specifies a maximum allowable time difference between injecting a computing request into the distributed computing cluster and receiving a response to the injected computing request from the distributed computing cluster.

4. The computing system of claim any one of claims 1 to 3 further including a monitor (119) configured to observe a rate of output of the distributed computing cluster and a processing latency of the distributed computing system.

5. The computing system of claim 4 wherein the observed processing latency is measured by determining a time difference between injecting a computing request into the distributed computing cluster and receiving a response to the injected computing request from the distributed computing cluster.

6. The computing system of claim 4 wherein requests injected into the distributed computing cluster and results output by the distributed computing cluster pass through the monitor.

7. The computing system of claim 4 wherein the computational throughput is observed by measuring, by the monitor, a rate of data elements outputted by the cluster.

8. The computing system of claim 4 wherein the controller is configured to process the observed computational throughput of the distributed computing cluster and the observed processing latency to adjust the value of the second parameter and the value of the first parameter.

9. The computing system of claim 8 wherein the controller repeatedly processes the observed computational throughput of the distributed computing cluster and the observed processing latency to adjust the value of the second parameter and the value of the first parameter while the distributed computing cluster is processing injected computing requests.

10. The computing system of claim 8 wherein the controller implements a feedback loop with the distributed computing cluster to process the observed computational throughput of the distributed computing cluster and the observed processing latency to adjust the value of the second parameter and the value of the first parameter.

11. The computing system of claim any one of claims 1 to 10 wherein the controller uses an optimization algorithm to control the value of the first and/or second parameter to maximize the computational throughput of the distributed computing cluster while complying with the at least one service level requirement.

12. The computing system of claim any one of claims 1 to 11 wherein the first parameter includes a delay parameter governing a minimum time interval between transmission of messages by the computing node over the interconnect network.

13. The computing system of claim 12 wherein the at least some computing nodes are configured to accumulate messages for transmission as a group of messages during the minimum time interval.

14. The computing system of claim any one of claims 1 to 13 wherein the second parameter includes a rate of input parameter configured to limit a rate of the injection of computing requests into the distributed computing cluster.

15. The computing system of claim any one of claims 1 to 14 wherein controlling the computational throughput of the distributed computing cluster while complying with the at least one service level requirements includes the maximizing a computational throughput of the distributed computing cluster without violating any of the at least one service level requirements.

16. The computing system of claim any one of claims 1 to 15 wherein the first parameter governs a number of messages accumulated between transmissions of messages by the computing node over the interconnect network.

17. The computing system of any one of claims 1 to 16 further including one or more interface components, each configured to receive input data, inject computing requests into the distributed computing cluster for processing the input data, receive processing results from the distributed computing cluster, and generate output data from the processing results, the interface component being configured to limit a rate of injection of computing requests into the distributed computing system according to the second parameter.

18. The computing system of claim 17 wherein at least some of the one or more interface components execute on a client system (106) separate from systems on which the distributed computing cluster executes.

19. The computing system of claim 17 wherein at least some of the received input data is associated with priority levels, and the one or more interface components are configured to manage injection of the input data into the distributed data cluster according to the priority levels.

20. A method including:

limiting injections of computing requests into a distributed computing cluster according to a second parameter, wherein the distributed computing system includes:

a plurality of computing nodes (105) interconnected by an interconnect network (107) over which the computing nodes of the plurality of computing nodes communicate with each other by passing messages, each of at least some of the computing nodes being configured with a first parameter governing transmissions of messages by the computing node over the interconnect network,

accumulating, at the at least some computing nodes, messages for transmission as a group of messages according to the first parameter,

receiving, at a controller (115), at least one predetermined service level requirement, and

controlling, using the controller, a value of the second parameter and a value of the first parameter to control a computational throughput of the distributed computing cluster while complying with the at least one service level requirement.

21. A computer-readable medium storing software in a non-transitory form, the software including instructions for causing a computing system to:

limit a rate of injection of computing requests into a distributed computing cluster according to a second parameter, wherein the distributed computing system includes:

a plurality of computing nodes (105) interconnected by an interconnect network (107) over which the computing nodes of the plurality of computing nodes communicate with each other by passing messages, each of at least some of the computing nodes being configured with a first parameter governing transmissions of messages by the computing node over the interconnect network,

accumulate, at the at least some computing nodes, messages for transmission as a group of messages according to the first parameter,

receive, at a controller (115), at least one predetermined service level requirement, and

control, using the controller, a value of the second parameter and a value of the first parameter to control a computational throughput of the distributed computing cluster while complying with the at least one service level requirement.

22. A computing system including:

means for limiting a rate of injection of computing requests into a distributed computing cluster according to a second parameter, wherein the distributed computing system includes:

a plurality of computing nodes (105) interconnected by an interconnect network (107) over which the computing nodes of the plurality of computing nodes communicate with each other by passing messages, each of at least some of the computing nodes being configured with a first parameter governing a transmissions of messages by the computing node over the interconnect network,

means for accumulating, at the at least some computing nodes, messages for transmission as a group of messages according to the first parameter,

means for receiving, at a controller (115), at least one predetermined service level agreement, and

means for controlling, using the controller, a value of the second parameter and a value of the first parameter to control a computational throughput of the distributed computing cluster while complying with the at least one service level requirement.