CLAIMS

What is claimed is:

1. A computing system (104) including:

a controller,

a monitor, and

a distributed computing cluster (112) that includes: computing nodes (105) and a rate limiter,

wherein the monitor is configured to observe both a rate of output of the distributed computing cluster and a processing latency of the distributed computing cluster,

wherein the computing nodes communicate with each other by passing messages over an interconnect network (107) that interconnects the computing nodes,

wherein at least some of the computing nodes are configured to accumulate messages for transmission as a group of messages according to a first parameter, wherein the first parameter governs transmission of messages by the computing node over the interconnect network, and

wherein the rate limiter is configured to limit injections of computing requests into the distributed computing cluster according to a second parameter, and

wherein the controller (110) is configured to receive at least one predetermined service level

requirement and to control values of the first and second parameters 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 at least one service level requirement 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.

3. The computing system of claim 1, 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.

4. The computing system of claim 1, wherein requests injected into the distributed computing cluster and results output by the distributed computing cluster pass through the monitor.

5. The computing system of claim 1, wherein the computational throughput is observed by having the monitor measure a rate at which the cluster outputs data elements.

6. The computing system of claim 1, 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.

7. The computing system of claim 6, 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.

8. The computing system of claim 6, 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.

9. The computing system of any one of claims 1 to 8, 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.

10. The computing system of any one of claims 1 to 9, 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.

11. The computing system of claim 10, wherein the at least some computing nodes are configured to accumulate messages for transmission as a group of messages during the minimum time interval.

12. The computing system of claim any one of claims 1 to 11, wherein the second parameter limits a rate at which

computing requests are injected into the distributed computing cluster.

13. The computing system of claim any one of claims 1 to 12, wherein the at least one predetermined service level requirement is one of a plurality of predetermined service level requirements, wherein controlling the computational throughput of the distributed computing cluster while complying with the at least one service level requirement~~s~~ includes maximizing a computational throughput of the distributed computing cluster without violating any of the service level requirements in the plurality of service level requirments.

14. The computing system of claim any one of claims 1 to 13, wherein the first parameter governs how many messages are to be accumulated between transmissions of messages by the computing node over the interconnect network.

15. The computing system of any one of claims 1 to 14, further including one or more interface components, each of which is configured to receive input data, to inject computing requests into the distributed computing cluster for processing the input data, to receive processing results from the distributed computing cluster, and to generate output data from the processing results, and to limit a rate of injection of computing requests into the distributed computing cluster according to the second parameter.

16. The computing system of claim 15, wherein at least some of the one or more interface components execute on a client

system (106) that is separate from systems on which the distributed computing cluster executes.

17. The computing system of claim 16, wherein at least some of the received input data is associated with priority levels and wherein 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.

18. A method including: limiting injections of computing requests into a distributed computing cluster that is part of a distributed computing system, wherein the distributed computing system includes: a controller and a monitor, wherein the distributed computing cluster includes a plurality of computing nodes (105) and a rate limiter, wherein the computing nodes communicate with each other by passing messages over an interconnect network (107) that interconnects the computing nodes, wherein limiting the injections of the computing requests includes: causing at least some of the computing nodes to accumulate messages for transmission as a group of messages according to a first parameter, wherein the first parameter governs transmission of messages by computing nodes over the interconnect network, receiving, at the controller (115), at least one predetermined service level requirement, using the monitor, observing both a rate of output of the distributed computing cluster and a processing latency of the distributed computing cluster, and using the controller, controlling a value of a 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.

19. A computer-readable medium storing software in a non-transitory form, the software including instructions for causing a distributed computing system, which includes a controller, a monitor, and a distributed computing cluster, to: limit a rate of injection of computing requests into the distributed computing cluster according to a second parameter, wherein the distributed computing cluster includes: a plurality of computing nodes (105) that communicate with each other by passing messages over an interconnect network (107) that interconnects the computing nodes, wherein each of at least some of the computing nodes are configured to accumulate messages for transmission as a group of messages according to a first parameter, wherein the first parameter governs transmission of messages by the computing node over the interconnect network, wherein the instructions comprise instructions to cause the controller (115) to receive, at least one predetermined service level requirement, instructions to cause the monitor to observe both a rate of output of the distributed computing cluster and a processing latency of the distributed computing cluster, and instructions to cause the controller 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.