

What is claimed is:

1. A method for processing data using a distributed data processing system including a plurality of processing nodes, the method including:
 - receiving input data at a distributed data processing system interface component in communication with the distributed data processing system;
 - providing the received input data to the distributed data processing system, wherein the distributed data processing system assigns time information associated with a first time interval of a plurality of time intervals to the input data;
 - generating state information within the distributed data processing system in response to receiving the input data, wherein the state information includes the indicator associated with the first time interval;
 - receiving and storing, at the distributed data processing system interface component, result data associated with the input data from the distributed data processing system, wherein the result data includes second time information associated with a second time interval later than the first time interval;
 - in response to determining, at a first processing node of the plurality of processing nodes that all data records associated with state information that includes a particular time or earlier have been stored at a first level of durability, updating an indicator with a value representative of the particular time;
 - providing the updated indicator to the distributed data processing system interface component; and
 - determining whether to release the result data from the distributed data processing system interface component based on a comparison, at the distributed data processing system interface component, of the updated indicator to the second time information associated with the second time interval included in the result data.

2. The method of claim 1 further comprising releasing the result data from the distributed data processing system interface component if the indicator corresponds to a time interval identical to or later than the second time interval.
3. The method of any one of claims 1 through 2 wherein the input data is stored in a buffer at the interface component and releasing result data includes removing input data from the buffer.
4. The method of any one of claims 1 through 3 wherein the first level of durability indicates that all state information is replicated in volatile memory associated with one or more processing nodes of the plurality of processing nodes of the distributed data processing system.
5. The method of any one of claims 1 through 4 wherein the first level of durability indicates that all state information is stored in persistent memory associated with one or more processing nodes of the plurality of processing nodes of the distributed data processing system.
6. The method of any one of claims 1 through 5 wherein the first level of durability indicates that all state information is stored in persistent memory outside of the distributed data processing system.
7. The method of any one of claims 1 through 6 further comprising processing the input data at a processing node of the plurality of processing nodes to generate the result data.
8. The method of any one of claims 1 through 7 wherein the input data includes data processing task, and the result data includes a result of the data processing task.
9. The method of any one of claims 1 through 8 wherein the input data is stored in a buffer at the interface component, the method further comprising receiving an indication that processing of the input data has failed from the first processing node, discarding the result, and re-providing the input data stored in the buffer to the distributed data processing system.

10. The method of claim 9 wherein the input data is associated with a unique identifier and the re-provided input data has the same unique identifier as the originally provided input data.
11. The method of any one of claims 1 through 10 ,wherein in response to determining, at the first processing node that all of one or more data records associated with respective one or more state information portions that include the particular time or earlier have been stored at a second level of durability, updating a second indicator with the value representative of the particular time.
12. The method of claim 11 wherein the second level of durability is more durable than the first level of durability.
13. The method of claim 11 wherein the second level of durability is less durable than the first level of durability.
14. Software stored in a non-transitory form on a computer-readable medium, for processing data using a distributed data processing system including a plurality of processing nodes, the software including instructions for causing a computing system to perform all of the steps of any one of claims 1-13.
15. An apparatus for processing data, the apparatus including:
a distributed data processing system including a plurality of processing nodes,
each processing node including at least one processor;
a distributed data processing system interface component in communication with the distributed data processing system; and
a communication medium connecting the plurality of processing nodes for sending and receiving information between processing nodes of the plurality of processing nodes;
wherein the apparatus is configured to perform the steps of the method as claimed in any one of claims 1 to 13
16. A computing system for processing data using a distributed data processing system including a plurality of processing nodes, the computing system including:

means for receiving input data at a distributed data processing system interface component in communication with the distributed data processing system;

means for providing the received input data to the distributed data processing system, wherein the distributed data processing system assigns time information associated with a first time interval of a plurality of time intervals to the input data;

means for generating state information within the distributed data processing system in response to receiving the input data, wherein the state information includes the indicator associated with the first time interval;

means for receiving and storing, at the distributed data processing system interface component, result data associated with the input data from the distributed data processing system, wherein the result data includes second time information associated with a second time interval later than the first time interval;

means for updating an indicator with a value representative of the particular time in response to determining, at a first processing node of the plurality of processing nodes that all data records associated with state information that includes a particular time or earlier have been stored at a first level of durability;

means for providing the updated indicator to the distributed data processing system interface component; and

means for determining whether to release the result data from the distributed data processing system interface component based on a comparison, at the distributed data processing system interface component, of the updated indicator to the second time information associated with the second time interval included in the result data.

17. The computing system of claim 16, further comprising means for releasing the result data from the distributed data processing system interface component if the indicator corresponds to a time interval identical to or later than the second time

interval.

18. The computing system of claim 17, wherein the input data is stored in a buffer at the interface component, and wherein the means for releasing the result data comprises means for removing the input data from the buffer.

19. The computing system of claim 16, wherein the first level of durability indicates one or more: a) that all state information is replicated in volatile memory associated with one or more processing nodes of the plurality of processing nodes of the distributed data processing system, b) that all the state information is stored in persistent memory associated with one or more processing nodes of the plurality of processing nodes of the distributed data processing system, or c) that all the state information is stored in persistent memory outside of the distributed data processing system.

20. The computing system of claim 16, wherein the input data is stored in a buffer at the interface component, and wherein the system further comprises means for receiving an indication that processing of the input data has failed from the first processing node, means for discarding the result, and means for re-providing the input data stored in the buffer to the distributed data processing system.

21. The computing system of claim 16, further comprising means for updating a second indicator with the value representative of the particular time in response to determining, at the first processing node, that all of one or more data records associated with respective one or more state information portions that include the particular time or earlier have been stored at a second level of durability.

22. The computing system of claim 21, wherein the second level of durability is one of: more durable than the first level of durability, less durable than the first level of durability.