CLAIMS

1. A method comprising:

receiving a request to transfer data from a first computing device in a network to a second computing device in the network, the request identifying:

the second computing device;

a volume of data to be transferred from the first computing device to the second computing device in accordance with the request; and

a deadline, the transfer of the data from the first computing device to the second computing device to be completed prior to the deadline;

responsive to receiving the request and based upon the request, executing an optimization process to compute a long-term schedule that covers a plurality of time units, the long-term schedule generated to facilitate completion of the transfer of the data from the first computing device to the second computing device prior to the deadline, the long-term schedule identifies, for a time unit in the plurality of time units, at least one path in the network over which the data is to be transferred from the first computing device to the second computing device;

the optimization process comprising:

receiving a map of the network, the map of the network comprising a plurality of nodes that are respectively representative of devices in the network, the map of the network further comprising a plurality of edges that are representative of links between respective devices in the network;

responsive to receipt of the map of the network, constructing a network graph that is representative of the network over the plurality of time units, the network graph comprises a plurality of instances of the map of the network that respectively correspond to the plurality of time units, the plurality of instances of the map of the network connected with one another by way of edges between nodes in the plurality of instances of the map of the network that represent common devices; and

generating the long-term schedule based upon the network graph;

based upon the long-term schedule, computing a short-term schedule, the short-term schedule computed to facilitate completion of the transfer of the volume of the data from the first computing device to the second computing device prior to the deadline, the short term schedule comprising a routing table for a network infrastructure device in the

network, the routing table identifies at least one device to which data received by the network infrastructure device is to be transferred; and

transmitting the routing table to the network infrastructure device.

2. The method of claim 1, further comprising, for the data transfer between the first computing device and the second computing device, constraining a number of paths represented in the network graph that can be used to transfer the data from the first computing device to the second computing device to a threshold number of potential paths, wherein generating the long-term schedule is based upon the threshold number of potential paths.

3. The method of claim 1, further comprising:

receiving a plurality of requests for a respective plurality of data transfers between various computing devices in the network, the plurality of requests identifying a respective plurality of deadlines, the respective plurality of deadlines being different; and

generating the long-term schedule based upon the plurality of requests, the long term schedule generated to facilitate completion of the data transfers between the various computing devices in the network prior to their respective plurality of deadlines.

4. The method of claim 1, wherein the optimization process is executed in parallel over multiple processor cores.

5. The method of claim 1, wherein the optimization process includes execution of a mixed packing and covering algorithm.

6. A controller computing device comprising:

a processor; and

a memory that comprises a plurality of components that are executed by the processor, the plurality of components comprising:

a receiver component that receives a plurality of data transfer requests having a respective plurality of specified deadlines, each data transfer request being a respective request to transfer data over a network that comprises a plurality of network infrastructure devices; and

a scheduler component that:

computes a long-term schedule based upon the plurality of data transfer requests, the long-term schedule covering a plurality of time units and defining flow of traffic through the network in the plurality of future time units, wherein the long-term schedule is computed by executing an optimization process, the optimization process comprising:

receiving a map of the network, the map of the network comprising a plurality of nodes that are respectively representative of devices in the network, the map of the network further comprising a plurality of edges that are representative of links between respective devices in the network;

responsive to receipt of the map of the network, constructing a network graph that is representative of the network over the plurality of time units, the network graph comprises a plurality of instances of the map of the network that respectively correspond to the plurality of time units, the plurality of instances of the map of the network connected with one another by way of edges between nodes in the plurality of instances of the map of the network that represent common devices; and

generating the long-term schedule based upon the network graph;

computes a short-term schedule based upon the long-term schedule, the short-term schedule covering fewer time units than the long-term schedule; and

causes respective routing tables to be transmitted to the plurality of network infrastructure devices, the plurality of network infrastructure devices routing data packets in the network based upon the routing tables.

7. The computing device of claim 6, the scheduler component executes a mixed packing and covering algorithm to compute the long-term schedule.

8. The computing device of claim 7, wherein the long-term schedule comprises a respective sub-schedule for each time unit, the respective sub-schedule defining flow of traffic through the network for its corresponding time unit.