**WHAT IS CLAIMED IS:**

1. A method comprising:

identifying, by a first node included in a mesh network, a plurality of potential parent nodes for the first node;

receiving, by the first node and for respective potential parent nodes from the plurality of potential parent nodes, respective accumulated uplink message success rates, wherein a respective accumulated uplink message success rate for a respective potential parent node indicates a probability of successfully transmitting messages from that respective potential parent node to a target destination;

receiving, by the first node from the respective potential parent nodes, respective accumulated downlink message success rates, wherein a respective accumulated downlink message success rate indicates a probability of successfully receiving, at that respective potential parent node, messages transmitted to that potential parent node from the target destination; and

selecting, by the first node based on at least one respective combined accumulated message success rate that is computed based on at least one of the respective accumulated uplink message success rates and at least one of the respective accumulated downlink message success rates, a first potential parent node from the plurality of potential parent nodes as a parent node.

2. The method of claim 1, further comprising:

computing, by the first node, respective combined accumulated message success rates for the respective potential parent nodes based on respective accumulated downlink message success rates and the respective accumulated uplink message success rates.

3. The method of claim 1, further comprising computing, by the first node, respective second accumulated uplink message success rates based on respective local uplink message success rates associated with transmitting messages from the first

node to the respective potential parent nodes and on the respective accumulated uplink message success rates received from the respective potential parent nodes.

4. The method of claim 3, wherein a first local uplink message success rate associated with the first potential parent node is greater than a second local uplink message success rate associated with a second potential parent node included in the plurality of potential parent nodes, and a first accumulated uplink message success rate associated with the first potential parent node is less than a second accumulated uplink message success rate associated with the second potential parent node, and selecting the parent node comprises selecting the second potential parent node, instead of the first potential parent node, to be the parent node for the first node.

5. The method of claim 3, wherein a first local uplink message success rate associated with the first potential parent node is greater than a second local uplink message success rate associated with a second potential parent node included in the plurality of potential parent nodes, and a first accumulated uplink message success rate associated with the first potential parent node is greater than a second accumulated uplink message success rate associated with the second potential parent node, and selecting the parent node comprises selecting the first potential parent node, instead of the second potential parent node, to be the parent node for the first node.

6. The method of claim 1, wherein selecting the parent node comprises:

determining, for the first potential parent node, a comparison message success rate from a respective accumulated uplink message success rate received by the first node from the first potential parent node; and

selecting, based on the first potential parent node having a highest comparison message success rate among the plurality of respective potential parent nodes, the first potential parent node as the parent node.

7. The method of claim 1, wherein selecting the parent node comprises filtering the plurality of potential parent nodes based on one or more message success rate criteria.

8. The method of claim 7, wherein selecting the parent node further comprises filtering the plurality of potential parent nodes based on one or more hierarchical secondary selection criteria to obtain the parent node for the first node.

9. One or more non-transitory computer-readable media storing instructions that, when executed by one or more processors, cause the one or more processors to perform the steps of:

identifying, by a first node included in a mesh network, a plurality of potential parent nodes for the first node;

receiving, by the first node and for respective potential parent nodes from the plurality of potential parent nodes, respective accumulated uplink message success rates, wherein a respective accumulated uplink message success rate for a respective potential parent node indicates a probability of successfully transmitting messages from that respective potential parent node to a target destination;

receiving, from the respective potential parent nodes, respective accumulated downlink message success rates, wherein a respective accumulated downlink message success rate for a respective potential parent node indicates a probability of successfully receiving messages transmitted to that respective potential parent node from the target destination; and

selecting, by the first node based at least on one respective combined accumulated message success rate that is computed based on at least one of the respective accumulated uplink message success rates and at least one of the respective accumulated downlink message success rates, a first potential parent node from the plurality of potential parent nodes as a parent node.

10. The one or more non-transitory computer-readable media of claim 9, wherein the steps further comprise

computing respective combined accumulated message success rates for the potential parent nodes based on the respective accumulated downlink

message success rates and the respective accumulated uplink message success rates.

11. The one or more non-transitory computer-readable media of claim 9, wherein the steps further comprise computing second accumulated message uplink success rates for respective potential parent nodes based on respective local uplink message success rates associated with transmitting messages from the first node to the potential parent nodes and on the respective accumulated message uplink success rates received from the respective potential parent nodes.

12. The one or more non-transitory computer-readable media of claim 11, wherein a first local uplink message success rate associated with the first potential parent node is greater than a second local uplink message success rate associated with a second potential parent node included in the plurality of potential parent nodes, and a first accumulated uplink message success rate associated with the first potential parent node is less than a second accumulated uplink message success rate associated with the second potential parent node, and selecting the parent node comprises selecting the second potential parent node, instead of the first potential parent node, to be the parent node for the first node.

13. The one or more non-transitory computer-readable media of claim 11, wherein a first local uplink message success rate associated with the first potential parent node is greater than a second local uplink message success rate associated with a second potential parent node included in the plurality of potential parent nodes, and a first accumulated uplink message success rate associated with the first potential parent node is greater than a second accumulated uplink message success rate associated with the second potential parent node, and selecting the parent node comprises selecting the first potential parent node, instead of the second potential parent node, to be the parent node for the first node.

14. The one or more non-transitory computer-readable media of claim 9, wherein selecting the parent node comprises:

determining, for the first potential parent node, a comparison message success rate from a respective accumulated uplink message success rate received from the first potential parent node; and

filtering the plurality of potential parent nodes based on one or more message success rate criteria associated with the comparison message success rate.

15. The one or more non-transitory computer-readable media of claim 14, wherein the one or more message success rate criteria includes a first criterion that specifies a minimum threshold value associated with comparison message success rates, and wherein filtering the plurality of potential parent nodes comprises determining, for the first potential parent node, whether the comparison message success rate exceeds the minimum threshold value.

16. The one or more non-transitory computer-readable media of claim 14, wherein the one or more message success rate criteria includes a first criterion that specifies a marginal threshold value associated with comparison message success rates, and wherein filtering the plurality of potential parent nodes comprises:

identifying the first potential parent node associated with a highest comparison message success rate; and

for potential parent nodes other than the first potential parent node, determining whether the comparison message success rate is within the marginal threshold value of the highest comparison message success rate.

17. A system comprising:

a memory that stores instructions, and

a processor that is coupled to the memory and, when executing the instructions, is configured to:

identify, by a first node included in a mesh network, a plurality of potential parent nodes for the first node;

receive, by the first node and for respective potential parent nodes from the plurality of potential parent nodes, respective accumulated uplink

message success rates, wherein a respective accumulated uplink message success rate for a respective potential parent node indicates a probability of successfully transmitting messages from that respective potential parent node to a target destination;

receive, by the first node from the respective potential parent nodes, respective accumulated downlink message success rates, wherein a respective accumulated downlink message success rate indicates a probability of successfully receiving, at that respective potential parent node, messages transmitted to that potential parent node from the target destination; and

select, by the first node based on at least one respective combined accumulated message success rate that is computed based on at least one of the respective accumulated uplink message success rates and at least one of the respective accumulated downlink message success rates, a first potential parent node from the plurality of potential parent nodes as a parent node.

18. The system of claim 17, wherein the processor, when executing the instructions, is further configured to compute, by the first node, respective second accumulated uplink message success rates based on respective local uplink message success rates associated with transmitting messages from the first node to the respective potential parent nodes and on the respective accumulated uplink message success rates received from the respective potential parent nodes.

19. The system of claim 17, wherein the processor is configured to select the parent node by:

determining, for the first potential parent node, a comparison message success rate from a respective accumulated uplink message success rate received by the first node from the first potential parent node; and

selecting, based on the first potential parent node having a highest comparison message success rate among the plurality of respective potential parent nodes, the first potential parent node as the parent node.

20. The system of claim 17, wherein the processor is configured to select the parent node by filtering the plurality of potential parent nodes based on one or more message success rate criteria.