WHAT IS CLAIMED IS:

1. A method comprising:

determining, at a first node in a mesh network, a first combined accumulated message success rate that is based on an accumulated uplink message success rate associated with transmitting messages from the first node to a target destination within the mesh network via an established parent node, a first offset value associated with the accumulated uplink message success rate, an accumulated downlink message success rate associated with receiving messages from the target destination via the established parent node, and a second offset value associated with the accumulated downlink message success rate;

determining, at the first node based on the first combined accumulated message success rate, that a search for a different parent node should be performed;

in response to determining that the search for the different parent node should be performed, at the first node, performing one or more communication operations to discover a plurality of nearby nodes;

computing, at the first node, one or more respective second combined accumulated message success rates that are based on respective accumulated uplink message success rates associated with transmitting messages from the first node to the target destination within the mesh network via a respective nearby node of the plurality of nearby nodes, respective first offset values associated with the respective accumulated uplink message success rates, respective accumulated downlink message success rates associated with receiving messages from the target destination via the respective nearby node, and respective second offset values associated with the respective accumulated downlink message success rates; and

based on the one or more respective second combined accumulated message success rates, selecting, at the first node, a new parent node from the plurality of nearby nodes or maintaining the established parent node.

2. The method of claim 1, wherein the accumulated uplink message success rate indicates a probability of successfully transmitting, from the first node, messages transmitted from the first node to the target destination via the established parent node.

3. The method of claim 2, wherein determining that the search for the different parent node should be performed comprises determining that the accumulated uplink message success rate is lower than a threshold amount.

4. The method of claim 1, wherein the accumulated uplink message success rate indicates a probability of successfully transmitting messages from the first node to the target destination via the established parent node and the accumulated downlink message success rate indicates a probability of successfully receiving, at the first node, messages transmitted to the first node from the target destination via the established parent node.

5. The method of claim 1, wherein determining that the search for the different parent node should be performed comprises determining that the first combined accumulated message success rate is lower than a threshold amount.

6. The method of claim 1, wherein the first combined accumulated message success rate is determined based on a product of the accumulated uplink message success rate, as offset by the first offset value associated with the accumulated uplink message success rate, and the accumulated downlink message success rate, as offset by the second offset value associated with the accumulated downlink message success rate.

7. The method of claim 1, wherein determining the first combined accumulated message success rate comprises:

scaling the accumulated uplink message success rate by a first scaling factor to determine a scaled accumulated uplink message success rate;

adding the first offset value to the scaled accumulated uplink message success rate to determine an offset accumulated uplink message success rate;

scaling the accumulated downlink message success rate by a second scaling factor to determine a scaled accumulated downlink message success rate;

adding the second offset value to the scaled accumulated downlink message success rate to determine an offset accumulated downlink message success rate; and

multiplying the offset accumulated uplink message success rate by the offset accumulated downlink message success rate to determine the first combined accumulated message success rate.

8. The method of claim 1, wherein selecting the new parent node from the plurality of nearby nodes or maintaining the established parent node comprises:

determining that a second combined accumulated message success rate associated with at least one nearby node is greater than the first combined accumulated message success rate associated with the established parent node; and

selecting the new parent node from the plurality of nearby nodes.

9. The method of claim 1, wherein selecting the new parent node from the plurality of nearby nodes or maintaining the established parent node comprises:

determining that no second combined accumulated message success rate associated with the plurality of nearby nodes is greater than the first combined accumulated message success rate associated with the established parent node; and

maintaining the established parent node.

10. The method of claim 1, wherein the established parent node is included in the plurality of nearby nodes, and wherein maintaining the established parent node comprises selecting the established parent node from the plurality of nearby nodes.

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

determining, at a first node, a first combined accumulated message success rate that is based on an accumulated uplink message success rate associated with transmitting messages from the first node to a target destination within a mesh network via an established parent node, a first offset value associated with the accumulated uplink message success rate, an accumulated downlink message success rate associated with receiving messages from the target destination via the established parent node, and a second offset value associated with the accumulated downlink message success rate;

determining, at the first node based on the first combined accumulated message success rates, that a search for a different parent node should be performed;

in response to determining that the search for the different parent node should be performed, at the first node, performing one or more communication operations to discover a plurality of nearby nodes;

computing, at the first node, one or more respective second combined accumulated message success rates that are based on respective accumulated uplink message success rates associated with transmitting messages from the first node to the target destination within the mesh network via a respective nearby node of the plurality of nearby nodes, respective first offset values associated with the respective accumulated uplink message success rates, respective accumulated downlink message success rates associated with receiving messages from the target destination via the respective nearby node, and respective second offset values associated with the respective accumulated downlink message success rates; and

based on the one or more respective second combined accumulated message success rates, selecting, at the first node, a new parent node from the plurality of nearby nodes or maintaining the established parent node.

12. The one or more non-transitory computer-readable media of claim 11, wherein the accumulated uplink message success rate indicates a probability of successfully transmitting, from the first node, messages transmitted from the first node to the target destination via the established parent node.

13. The one or more non-transitory computer-readable media of claim 12, wherein determining that the search for the different parent node should be performed comprises determining that the accumulated uplink message success rate is lower than a threshold amount.

14. The one or more non-transitory computer-readable media of claim 11, wherein the accumulated uplink message success rate that indicates a probability of successfully transmitting messages from the first node to the target destination via the established parent node and the accumulated downlink message success rate indicates a probability of successfully receiving, at the first node, messages transmitted to the first node from the target destination via the established parent node.

15. The one or more non-transitory computer-readable media of claim 11, wherein determining that the search for the different parent node should be performed comprises determining that the first combined accumulated message success rate is lower than a threshold amount.

16. The one or more non-transitory computer-readable media of claim 11, wherein selecting the new parent node from the plurality of nearby nodes or maintaining the established parent node comprises:

determining respective comparison message success rates from the one or more respective second combined accumulated message success rates for the plurality of nearby nodes;

determining that a comparison message success rate associated with at least one nearby node is greater than a comparison message success rate associated with the established parent node; and

selecting the new parent node from the plurality of nearby nodes.

17. The one or more non-transitory computer-readable media of claim 11 wherein the established parent node is included in the plurality of nearby nodes, and wherein maintaining the established parent node comprises selecting the established parent node from the plurality of nearby nodes.

18. The one or more non-transitory computer-readable media of claim 11 wherein selecting the new parent node from the plurality of nearby nodes or maintaining the established parent node comprises selecting the new parent node and establishing a communication link with the new parent node.

19. The one or more non-transitory computer-readable media of claim 11 wherein selecting the new parent node from the plurality of nearby nodes or maintaining the established parent node comprises determining, for each nearby node, whether the nearby node satisfies one or more message success rate criteria.

20. 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:

determine, at a first node, a first combined accumulated message success rate that is based on an accumulated uplink message success rate associated with transmitting messages from the first node to a target destination within a mesh network via an established parent node, a first offset value associated with the accumulated uplink message success rate, an accumulated downlink message success rate associated with receiving messages from the target destination via the established parent node, and a second offset value associated with the accumulated downlink message success rate;

determine, at the first node, based on the first combined accumulated message success rates, that a search for a different parent node should be performed;

in response to determining that the search for the different parent node should be performed, at the first node, perform one or more communication operations to discover a plurality of nearby nodes;

compute, at the first node, one or more respective second combined accumulated message success rates that are based on respective accumulated uplink message success rates associated with transmitting messages from the first node to the target destination within the mesh network via a respective nearby node of the plurality of nearby nodes, respective first offset values associated with the respective accumulated uplink message success rates, respective accumulated downlink message success rates associated with receiving messages from the target destination via the nearby node, and respective second offset values associated with the respective accumulated downlink message success rates; and

based on the one or more respective second combined accumulated message success rates, select, at the first node, a new parent node from the plurality of nearby nodes or maintaining the established parent node.