Distributed averaging (also known as average consensus) is an algorithm that builds on neighbor to neighbor interactions with the ultimate goal of convergence to the average of all initial node values or to some value close to that average. We analyze in this paper the performance of distributed averaging algorithms where the information exchanged between neighboring agents is subject to deterministic uniform quantization (i.e., real values sent by nodes to their neighbors are truncated). With such quantization, convergence to the precise average cannot be achieved in general, but the convergence would be to some value close to it, called quantized consensus. Using Lyapunov stability analysis, we characterize the convergence properties of the resulting nonlinear quantized system. We show that in finite time and depending on initial conditions, the algorithm on any static and connected graph will either cause all agents to reach a quantized consensus where the consensus value is the largest quantized value not greater than the average of their initial values, or will lead all variables to oscillate in a small neighborhood around the average. In the latter case, we identify tight bounds for the size of the neighborhood and we further show that the error can be made arbitrarily small by adjusting the algorithm's parameters in a distributed manner.