CLAIMS

1. A method for compressing audio data comprising: obtaining a sequence of digitized samples of an audio signal; performing a transform using the sequence of digitized samples, to generate a plurality of spectral lines; obtaining a group of spectral lines from the plurality of spectral lines; quantizing the group of spectral lines to generate a group of quantized values; wherein quantizing the group of spectral lines to generate the group of quantized values comprises: performing a specialized rounding operation on a spectral line selected from the group of spectral lines, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines; and using the specialized rounding operation to force a group parity value, computed for the group of quantized values, to a predetermined parity value; and outputting one or more data frames based on the group of quantized values; compressing the one or more data frames to generate compressed digital data; and encoding the compressed digital data to generate one or more channel bits.

2. The method of claim 1, wherein: the specialized rounding operation is performed on a pre-rounding value associated with the selected spectral line, the pre-rounding value comprises a floating-point value or a fixed-point value, and the group of quantized values comprise a group of integers or fixed-point values.

3. The method of claim 2, wherein the specialized rounding operation flips a rounding direction used for rounding the pre-rounding value associated with the selected spectral line, to force the group parity value computed for the group of quantized values to the predetermined parity value.

4. The method of claim 3, wherein the selected spectral line is selected for being associated with a pre-rounding value having a minimal distance to a midpoint between two nearest possible quantized values, as compared to other spectral lines in the group of spectral lines.

5. The method of claim 1, wherein in a tie between a first spectral line associated with a first pre-rounding value having a first distance to a midpoint between two nearest possible quantized values and a second spectral line associated with a second pre-rounding value having a second distance to a midpoint between two nearest possible quantized values, the first distance being equal to the second distance, the first spectral line is selected for being associated with a higher frequency bin of the transform than the second spectral line.

6. The method of claim 1, wherein the group of spectral lines includes a first spectral line associated with a first frequency bin and a first pre-rounding value and a second spectral line associated with a second frequency bin and a second pre-rounding value, wherein the first frequency bin corresponds to a higher frequency bin of the transform than the second frequency bin, wherein the first pre-rounding value corresponds to a first distance between two nearest possible quantized values, and the second pre-rounding value corresponds to a second distance between two nearest possible quantized values, the second distance being less than the first distance, and wherein the first spectral line is selected over the second spectral line.

7. The method of claim 1, wherein the one or more data frames comprise a group of codewords based on the group of quantized values.

8. The method of claim 7, wherein the group of codewords is generated from the group of quantized values using arithmetic encoding.

9. The method of claim 7, wherein the one or more data frames further comprise, for at least one quantized value in the group of quantized values, a rounding residual value.

10. The method of claim 9, wherein the one or more data frames further comprise, for the at least one quantized value in the group of quantized values, a parity residual value.

11. The method of claim 10, wherein the rounding residual value and the parity residual value are inserted in place of padding bits in the one or more data frames.

12. The method of claim 1, wherein the specialized rounding operation is used to force a sequence of groups of quantized values, quantized from a sequence of groups of spectral lines from the plurality of spectral lines, to have a sequence of predetermined parity values.

13. The method of claim 12, wherein the sequence of predetermined parity values is used as a watermark.

14. The method of claim 13, wherein the watermark signifies use of the specialized rounding operation.

15. The method of claim 14, wherein the watermark signifies presence of one or more parity residual values in the one or more data frames.

16. The method of claim 13, wherein the watermark is associated with a specific provider of a device implementing the method for compressing audio data.

17. The method of claim 1, wherein the one or more data frames maintain compatibility with an existing standard for audio data compression.

18. An encoder for compressing audio data comprising: a transform compute device configured to receive a sequence of digitized samples of an audio signal and compute a transform using the sequence of digitized samples of the audio signal, to generate a plurality of spectral lines; a quantizer configured to obtain a group of spectral lines from the plurality of spectral lines and quantize the group of spectral lines to generate a group of quantized values, wherein the quantizer is configured to quantize the group of spectral lines to generate the group of quantized values by performing a specialized rounding operation on a spectral line selected from the group of spectral lines, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines, and wherein the quantizer is configured to use the specialized rounding operation to force a group parity value, computed for the group of quantized values, to a predetermined parity value, and wherein the quantizer is further configured to output one or more data frames based on the group of quantized values; a compression device configured to compress the one or more data frames to generate compressed digital data; and an encoder configured to encode the compressed digital data to generate one or more channel bits.

19. The encoder of claim 18, wherein: the quantizer is configured to perform the specialized rounding operation on a pre-rounding value associated with the selected spectral line, the pre-rounding value comprises a floating-point value or a fixed-point value, and the group of quantized values comprise a group of integers or fixed-point values.

20. The encoder of claim 19, wherein the specialized rounding operation flips a rounding direction used for rounding the pre-rounding value associated with the selected spectral line, to force the group parity value computed for the group of quantized values to the predetermined parity value.

21. The encoder of claim 20, wherein the selected spectral line is selected for being associated with a pre-rounding value having a minimal distance to a midpoint between two nearest possible quantized values, as compared to other spectral lines in the group of spectral lines.

22. The encoder of claim 18, wherein in a tie between a first spectral line associated with a first pre-rounding value having a first distance to a midpoint between two nearest possible quantized values and a second spectral line associated with a second pre-rounding value having a second distance to a midpoint between two nearest possible quantized values, the first distance being equal to the second distance, the first spectral line is selected for being associated with a higher frequency bin of the transform than the second spectral line.

23. The encoder of claim 18, wherein the group of spectral lines includes a first spectral line associated with a first frequency bin and a first pre-rounding value and a second spectral line associated with a second frequency bin and a second pre-rounding value, wherein the first frequency bin corresponds to a higher frequency bin of the transform than the second frequency bin, wherein the first pre-rounding value corresponds to a first distance between two nearest possible quantized values, and the second pre-rounding value corresponds to a second distance between two nearest possible quantized values, the second distance being less than the first distance, and wherein the first spectral line is selected over the second spectral line.

24. The encoder of claim 18, wherein the one or more data frames comprise a group of codewords based on the group of quantized values.

25. The encoder of claim 24, wherein the quantizer is configured to generate the group of codewords from the group of quantized values using arithmetic encoding.

26. The encoder of claim 24, wherein the one or more data frames further comprise, for at least one quantized value in the group of quantized values, a rounding residual value.

27. The encoder of claim 26, wherein the one or more data frames further comprise, for the at least one quantized value in the group of quantized values, a parity residual value.

28. The encoder of claim 27, wherein the quantizer is configured to insert the rounding residual value and the parity residual value in place of padding bits in the one or more data frames.

29. The encoder of claim 19, wherein the quantizer is configured to use the specialized rounding operation to force a sequence of groups of quantized values, quantized from a sequence of groups of spectral lines from the plurality of spectral lines, to have a sequence of predetermined parity values.

30. The encoder of claim 29, wherein the sequence of predetermined parity values is used as a watermark.

31. The encoder of claim 30, wherein the watermark signifies use of the specialized rounding operation.

32. The encoder of claim 31, wherein the watermark signifies presence of one or more parity residual values in the one or more data frames.

33. The encoder of claim 30, wherein the watermark is associated with a specific provider of a device implementing a method for compressing audio data.

34. The encoder of claim 18, wherein the one or more data frames maintain compatibility with an existing standard for audio data compression.

35. A non-transitory computer-readable medium storing instructions therein for execution by one or more processing units, comprising instructions to: obtain a sequence of digitized samples of an audio signal; perform a transform using the sequence of digitized samples, to generate a plurality of spectral lines; obtain a group of spectral lines from the plurality of spectral lines; quantize the group of spectral lines to generate a group of quantized values; wherein the instructions to quantize the group of spectral lines to generate the group of quantized values comprise instructions to: perform a specialized rounding operation on a spectral line selected from the group of spectral lines, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines; and use the specialized rounding operation to force a group parity value, computed for the group of quantized values, to a predetermined parity value; and output one or more data frames based on the group of quantized values; compress the one or more data frames to generate compressed digital data; and encode the compressed digital data to generate one or more channel bits.

36. The non-transitory computer-readable medium of claim 35, wherein: the specialized rounding operation is performed on a pre-rounding value associated with the selected spectral line, the pre-rounding value comprises a floating-point value or a fixed-point value, and the group of quantized values comprise a group of integers or fixed-point values.

37. The non-transitory computer-readable medium of claim 36, wherein the specialized rounding operation flips a rounding direction used for rounding the pre-rounding value associated with the selected spectral line, to force the group parity value computed for the group of quantized values to the predetermined parity value.

38. The non-transitory computer-readable medium of claim 37, wherein the selected spectral line is selected for being associated with a pre-rounding value having a minimal distance to a midpoint between two nearest possible quantized values, as compared to other spectral lines in the group of spectral lines.

39. The non-transitory computer-readable medium of claim 36, wherein in a tie between a first spectral line associated with a first pre-rounding value having a first distance to a midpoint between two nearest possible quantized values and a second spectral line associated with a second pre-rounding value having a second distance to a midpoint between two nearest possible quantized values, the first distance being equal to the second distance, the first spectral line is selected for being associated with a higher frequency bin of the transform than the second spectral line.

40. The non-transitory computer-readable medium of claim 36, wherein the group of spectral lines includes a first spectral line associated with a first frequency bin and a first pre-rounding

value and a second spectral line associated with a second frequency bin and a second pre-rounding value, wherein the first frequency bin corresponds to a higher frequency bin of the transform than the second frequency bin, wherein the first pre-rounding value corresponds to a first distance between two nearest possible quantized values, and the second pre-rounding value corresponds to a second distance between two nearest possible quantized values, the second distance being less than the first distance, and wherein the first spectral line is selected over the second spectral line.

41. The non-transitory computer-readable medium of claim 36, wherein the one or more data frames comprise a group of codewords based on the group of quantized values.

42. The non-transitory computer-readable medium of claim 41, wherein the group of codewords is generated from the group of quantized values using arithmetic encoding.

43. The non-transitory computer-readable medium of claim 41, wherein the one or more data frames further comprise, for at least one quantized value in the group of quantized values, a rounding residual value.

44. The non-transitory computer-readable medium of claim 43, wherein the one or more data frames further comprise, for the at least one quantized value in the group of quantized values, a parity residual value.

45. The non-transitory computer-readable medium of claim 44, wherein the rounding residual value and the parity residual value are inserted in place of padding bits in the one or more data frames.

46. The non-transitory computer-readable medium of claim 35, wherein the specialized rounding operation is used to force a sequence of groups of quantized values, quantized from a sequence of groups of spectral lines from the plurality of spectral lines, to have a sequence of predetermined parity values.

47. The non-transitory computer-readable medium of claim 46, wherein the sequence of predetermined parity values is used as a watermark.

48. The non-transitory computer-readable medium of claim 47, wherein the watermark signifies use of the specialized rounding operation.

49. The non-transitory computer-readable medium of claim 48, wherein the watermark signifies presence of one or more parity residual values in the one or more data frames.

50. The non-transitory computer-readable medium of claim 47, wherein the watermark is associated with a specific provider of a device implementing a method for compressing audio data.

51. The non-transitory computer-readable medium of claim 35, wherein the one or more data frames maintain compatibility with an existing standard for audio data compression.

52. A system for compressing audio data comprising; means for obtaining a sequence of digitized samples of an audio signal; means for performing a transform using the sequence of digitized samples, to generate a plurality of spectral lines; means for obtaining a group of spectral lines from the plurality of spectral lines; means for quantizing the group of spectral

lines to generate a group of quantized values; wherein the means for quantizing the group of spectral lines to generate the group of quantized values comprises: means for performing a specialized rounding operation on a spectral line selected from the group of spectral lines, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines; means for using the specialized rounding operation to force a group parity value, computed for the group of quantized values, to a predetermined parity value; and means for outputting one or more data frames based on the group of quantized values; means for compressing the one or more data frames to generate compressed digital data; and means for encoding the compressed digital data to generate one or more channel bits.

53. The system of claim 52, wherein: the specialized rounding operation is performed on a pre-rounding value associated with the selected spectral line, the pre-rounding value comprises a floating-point value or a fixed-point value, and the group of quantized values comprise a group of integers or fixed-point values.

54. The system of claim 53, wherein the specialized rounding operation flips a rounding direction used for rounding the pre-rounding value associated with the selected spectral line, to force the group parity value computed for the group of quantized values to the predetermined parity value.

55. The system of claim 54, wherein the selected spectral line is selected for being associated with a pre-rounding value having a minimal distance to a midpoint between two nearest possible quantized values, as compared to other spectral lines in the group of spectral lines.

56. The system of claim 52, wherein in a tie between a first spectral line associated with a first pre-rounding value having a first distance to a midpoint between two nearest possible quantized values and a second spectral line associated with a second pre-rounding value having a second distance to a midpoint between two nearest possible quantized values, the first distance being equal to the second distance, the first spectral line is selected for being associated with a higher frequency bin of the transform than the second spectral line.