

WHAT IS CLAIMED IS:

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; 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.
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 4, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines.

6. The method of claim 5, 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.

7. The method of claim 5,
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.

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

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

10. The method of claim 8, 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.

11. The method of claim 10, 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.

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

13. 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.

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

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

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

17. The method of claim 14, wherein the watermark is associated with a specific provider of a device implementing the method for compressing audio data.

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

19. A method for de-compressing audio data comprising:
obtaining one or more data frames;
obtaining a group of quantized values based on the one or more data frames, wherein the group of quantized values originates from a compression-side quantization process involving a specialized rounding operation performed on a spectral line to force a parity value, computed for the group of quantized values, to a predetermined parity value;

computing a receive-side parity value for the group of quantized values;
comparing the computed receive-side parity value to the predetermined parity value for the group of quantized values;
performing a bit error operation to detect or correct at least one bit error in the one or more data frames, in response to detecting a difference between the computed receive-side parity value and the predetermined parity value for the group of quantized values;
estimating a group of spectral lines based on the group of quantized values, taking into account detection or correction of the at least one bit error in the one or more data frames;
performing an inverse transform using a plurality of spectral lines, including the group of spectral lines, to generate a sequence of digitized samples; and
outputting the sequence of digitized samples as a digital representation of an audio signal.

20. The method of claim 19, wherein:
the one or more data frames comprise a group of codewords, and
the bit error operation is performed to detect or correct at least one bit error in the group of codeword.

21. The method of claim 20, wherein the at least one bit error in the group of codewords is corrected by:
obtaining multiple transmissions of the group of codewords;
generating a plurality of reconstructed versions the group of codewords; and
selecting one reconstructed version of the group of codewords from the plurality of reconstructed versions of the group of codewords, based on a match between (1) the computed receive-side parity value associated with the one reconstructed version of the group of codewords and (2) the predetermined parity value.

22. The method of claim 21,
wherein a weak bit mask indicating positions of possible bit errors is generated by comparing the multiple transmissions of the group of codewords; and

wherein each of the plurality of reconstructed versions of the group of codewords is reconstructed by changing a bit at one of the bit positions indicated by the weak bit mask.

23. The method of claim 22, wherein the multiple transmissions of the group of codewords comprise (1) an original transmission of the group of codewords and (2) one or more re-transmissions of the group of codewords.

24. The method of claim 23, wherein each of the one or more re-transmissions of the group of codewords is triggered by a failed CRC associated with a previous transmission of the group of codewords.

25. The method of claim 20, wherein the group of quantized values are generated from the group of codewords using arithmetic decoding.

26. The method of claim 20, 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 method 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 method of claim 27, wherein the rounding residual value and the parity residual value are extracted from positions of padding bits in the one or more data frames.

29. The method of claim 27, wherein a spectral line from the group of spectral lines is estimated with increased resolution by taking into account the rounding residual value and the parity residual value.

30. The method of claim 29,
wherein the rounding residual value indicates a first estimated range of values for the spectral line, and the parity residual value indicates a second estimated range of values for the spectral line adjacent to the first estimated range, and
wherein the spectral line is estimated based on the second estimated range of values.

31. The method of claim 19, wherein a sequence of parity values are computed from a sequence of groups of quantized values obtained based on the one or more data frames.

32. The method of claim 31, wherein the sequence of predetermined parity values is used as a watermark.

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

34. 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; and
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 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.

35. The encoder of claim 34, 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.

36. The encoder of claim 35, 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.

37. The encoder of claim 36, 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.

38. The encoder of claim 37, wherein the selected spectral line is selected based on a selection bias favoring higher frequency spectral lines.

39. The encoder of claim 38, 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 encoder of claim 38,
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 encoder of claim 34, wherein the one or more data frames comprise a group of codewords based on the group of quantized values.

42. The encoder of claim 41, wherein the quantizer is configured to generate the group of codewords from the group of quantized values using arithmetic encoding.

43. The encoder 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 encoder 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 encoder of claim 44, 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.

46. The encoder of claim 35, 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.

47. A decoder for de-compressing audio data comprising:
a de-quantizer configured to receive one or more data frames and obtain a group of quantized values based on the one or more data frames, wherein the group of quantized values originates from a compression-side quantization process involving a

specialized rounding operation performed on a spectral line to force a parity value, computed for the group of quantized values, to a predetermined parity value,

wherein the de-quantizer is further configured to compute a receive-side parity value for the group of quantized values,

wherein the de-quantizer is further configured to compare the computed receive-side parity value to the predetermined parity value for the group of quantized values,

wherein the de-quantizer is further configured to perform a bit error operation to detect or correct at least one bit error in the one or more data frames, in response to detecting a difference between the computed receive-side parity value and the predetermined parity value for the group of quantized values, and

wherein the de-quantizer is further configured to estimate a group of spectral lines based on the group of quantized values, taking into account detection or correction of the at least one bit error in the one or more data frames; and

an inverse transform compute device configured to perform an inverse transform using a plurality of spectral lines, including the group of spectral lines, to generate and output a sequence of digitized samples, as a digital representation of an audio signal.

48. The decoder of claim 47, wherein:
the one or more data frames comprise a group of codewords, and
the bit error operation is performed to detect or correct at least one bit error in the group of codewords.

49. The decoder of claim 48, wherein the de-quantizer is configured to correct the at least one bit error in the group of codewords by:
obtaining multiple transmissions of the group of codewords;
generating a plurality of reconstructed versions the group of codewords; and
selecting one reconstructed version of the group of codewords from the plurality of reconstructed versions of the group of codewords, based on a match between (1) the computed receive-side parity value associated with the one reconstructed version of the group of codewords and (2) the predetermined parity value.

50. The decoder of claim 49,
wherein the de-quantizer is configured to generate a weak bit mask indicating positions of possible bit errors, by comparing the multiple transmissions of the group of codewords; and

wherein the de-quantizer is configured to reconstruct each of the plurality of reconstructed versions of the group of codewords by changing a bit at one of the bit positions indicated by the weak bit mask.

51. The decoder of claim 48, wherein the de-quantizer is configured to generate the group of integers from the group of codewords using arithmetic decoding.

52. The decoder of claim 48, 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.

53. The decoder of claim 52, 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.

54. The decoder of claim 53, wherein the de-quantizer is configured to extract the rounding residual value and the parity residual value from positions of padding bits in the one or more data frames.

55. The decoder of claim 53, wherein the de-quantizer is configured to estimate a spectral line from the group of spectral lines with increased resolution by taking into account the rounding residual value and the parity residual value.

56. The decoder of claim 55,
wherein the rounding residual value indicates a first estimated range of values for the spectral line, and the parity residual value indicates a second estimated range of values for the spectral line adjacent to the first estimated range, and
wherein the de-quantizer is configured to estimate the spectral line based on the second estimated range of values.

57. 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; 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.

58. A non-transitory computer-readable medium storing instructions therein for execution by one or more processing units, comprising instructions to:

- obtain one or more data frames;
- obtain a group of quantized values based on the one or more data frames, wherein the group of quantized values originates from a compression-side quantization process involving a specialized rounding operation performed on a spectral line to force a parity value, computed for the group of quantized values, to a predetermined parity value;
- compute a receive-side parity value for the group of quantized values;
- compare the computed receive-side parity value to the predetermined parity value for the group of quantized values;
- perform a bit error operation to detect or correct at least one bit error in the one or more data frames, in response to detecting a difference between the computed receive-side parity value and the predetermined parity value for the group of quantized values;
- estimate a group of spectral lines based on the group of quantized values, taking into account detection or correction of the at least one bit error in the one or more data frames;

perform an inverse transform using a plurality of spectral lines, including the group of spectral lines, to generate a sequence of digitized samples; and
output the sequence of digitized samples as a digital representation of an audio signal.

59. 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; and
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.

60. A system for de-compressing audio data comprising:
means for obtaining one or more data frames;
means for obtaining a group of quantized values based on the one or more data frames, wherein the group of quantized values originates from a compression-side quantization process involving a specialized rounding operation performed on a spectral line to force a parity value, computed for the group of quantized values, to a predetermined parity value;
means for computing a receive-side parity value for the group of quantized values;
means for comparing the computed receive-side parity value to the predetermined parity value for the group of quantized values;

means for performing a bit error operation to detect or correct at least one bit error in the one or more data frames, in response to detecting a difference between the computed receive-side parity value and the predetermined parity value for the group of quantized values;

means for estimating a group of spectral lines based on the group of quantized values, taking into account detection or correction of the at least one bit error in the one or more data frames;

means for performing an inverse transform using a plurality of spectral lines, including the group of spectral lines, to generate a sequence of digitized samples; and

means for outputting the sequence of digitized samples as a digital representation of an audio signal.