CLAIMS

1. A method for encoding an N-channel audio program, wherein the program is specified over a time interval, the time interval includes a subinterval from a time t1 to a time t2, and a time-varying mix, A(t), of N encoded signal channels to M output channels has been specified over the time interval, where M is less than or equal to N, said method including steps of:

determining a first cascade of N×N primitive matrices which, when applied to samples of the N encoded signal channels, implements a first mix of audio content of the N encoded signal channels to the M output channels, wherein the first mix is consistent with the time-varying mix, A(t), in the sense that the first mix is at least substantially equal to A(t1), and wherein an N×N primitive matrix is defined as a matrix in which N-1 rows contain off-diagonal elements equal to zero and on-diagonal elements with an absolute value of 1;

determining interpolation values which, with the first cascade of primitive matrices and an interpolation function defined over the subinterval, are indicative of a sequence of cascades of N×N updated primitive matrices, such that each of the cascades of updated primitive matrices, when applied to samples of the N encoded signal channels, implements an updated mix, associated with a different time in the subinterval, of the N encoded signal channels to the M output channels, wherein each said updated mix is consistent with the time-varying mix, A(t); and

generating an encoded bitstream which is indicative of encoded audio content, the interpolation values, and the first cascade of primitive matrices.

2. The method of claim 1, wherein each of the primitive matrices is a unit primitive matrix, the method further including:

a step of generating the encoded audio content by performing matrix operations on samples of the program’s N channels, including by applying a sequence of matrix cascades to the samples, wherein each matrix cascade in the sequence is a cascade of primitive matrices, and the sequence of matrix cascades includes a first inverse matrix cascade which is a cascade of inverses of the primitive matrices of the first cascade; or

a step of generating the encoded audio content by performing matrix operations on samples of the program's N channels, including by applying a sequence of matrix cascades to the samples, wherein each matrix cascade in the sequence is a cascade of primitive matrices, and each matrix cascade in the sequence is the inverse of a corresponding cascade of the cascades of N×N updated primitive matrices, and N = M, so that the M output channels are the same as the N channels of the program recovered losslessly.

3. The method of claim 1, wherein the first cascade of primitive matrices implements a seed primitive matrix, and the interpolation values are indicative of a seed delta matrix for the seed primitive matrix.

4. The method of claim 2, wherein a time-varying downmix, $A_2(t)$, of audio content or encoded content of the program to M1 speaker channels has also been specified over the time interval, where M1 is an integer less than M, and the method also includes a step of:

determining a second cascade of M1×M1 primitive matrices which, when applied to samples of M1channels of the encoded audio content at each time instant t in the interval implements a downmix of the N-channel audio program to the M1 speaker channels, wherein the downmix is consistent with the time-varying mix, $A_2(t)$.

5. A method for recovery of M channels of an N-channel audio program, wherein the program is specified over a time interval, the time interval includes a subinterval from a time t1 to a time t2, and a time-varying mix, A(t), of N encoded signal channels to M output channels has been specified over the time interval, said method including steps of:

obtaining an encoded bitstream which is indicative of encoded audio content, interpolation values, and a first cascade of N×N primitive matrices, wherein an N×N primitive matrix is defined as a matrix in which N-1 rows contain off-diagonal elements equal to zero and on-diagonal elements with an absolute value of 1; and

performing interpolation to determine a sequence of cascades of N×N updated primitive matrices, from the interpolation values, the first cascade of primitive matrices, and an interpolation function over the subinterval, wherein

the first cascade of N×N primitive matrices, when applied to samples of N encoded signal channels of the encoded audio content, implements a first mix of audio content of the N encoded signal channels to the M output channels, wherein the first mix is consistent with the

time-varying mix, A(t), in the sense that the first mix is at least substantially equal to A(t1), and

the interpolation values, with the first cascade of primitive matrices, and the interpolation function, are indicative of a sequence of cascades of N×N updated primitive matrices, such that each of the cascades of updated primitive matrices, when applied to samples of the N encoded signal channels of the encoded audio content, implements an updated mix, associated with a different time in the subinterval, of the N encoded signal channels to the M output channels, wherein each said updated mix is consistent with the time-varying mix, A(t).

6. The method of claim 5, wherein each of the primitive matrices is a unit primitive matrix.

7. The method of claim 6, wherein:

the encoded audio content has been generated by performing matrix operations on samples of the program's N channels, including by applying a sequence of matrix cascades to the samples, wherein each matrix cascade in the sequence is a cascade of primitive matrices, and the sequence of matrix cascades includes a first inverse matrix cascade which is a cascade of inverses of the primitive matrices of the first cascade; or

the encoded audio content has been generated by performing matrix operations on samples of the program's N channels, including by applying a sequence of matrix cascades to the samples, wherein each matrix cascade in the sequence is a cascade of primitive matrices, and each matrix cascade in the sequence is the inverse of a corresponding cascade of the cascades of N×N updated primitive matrices, and N = M, so that the M output channels are the same as the N channels of the program recovered losslessly.

8. The method of claim 7, wherein a time-varying downmix, $A_2(t)$, of audio content or encoded content of the program to M1 speaker channels has also been specified over the time interval, where M1 is an integer less than N, and the method also includes steps of:

receiving a second cascade of M1×M1 primitive matrices; and

applying the second cascade of M1×M1 primitive matrices to samples of M1 channels of the encoded audio content at each time instant t in the interval to implement a downmix of the N-channel audio program to the M1 speaker channels, wherein the downmix is consistent with the time-varying mix, $A_2(t)$.

9. The method of claim 5, wherein:

the encoded bitstream is also indicative of the interpolation function; or

the first cascade of primitive matrices implements a seed primitive matrix, and the interpolation values are indicative of a seed delta matrix for the seed primitive matrix.

10. The method of claim 5, said method also including a step of:

applying at least one of the cascades of updated N×N primitive matrices to samples of the encoded audio content, including by applying a seed primitive matrix and a seed delta matrix separately to the samples of the encoded audio content to generate transformed samples, and linearly combining the transformed samples in accordance with the interpolation function, thereby generating recovered samples indicative of samples of the M channels of the N-channel audio program.

11. The method of claim 5, wherein:

the interpolation function is substantially constant over some intervals of the encoded bitstream, and each most recently updated one of the cascades of N×N updated primitive matrices is updated by interpolation only during an interval of the encoded bitstream in which the interpolation function is not substantially constant; or

wherein the interpolation values include normalized delta values representable with Y bits, an indication of this number of bits of precision, and precision values, where the normalized delta values are indicative of normalized versions of delta values, the delta values are indicative of rates of change of coefficients of the primitive matrices, and the precision values indicative an increase in precision required to represent the delta values relative to the precision required to represent the coefficients of the primitive matrices.

12. The method of claim 11, wherein the delta values are derived by scaling the normalized delta values by a scale factor that is dependent on the resolution of the coefficients of the primitive matrices and the precision values.

13. The method of claim 5, also including steps of:

extracting a check word from the encoded bitstream, and verifying whether channels of a segment of the audio program have been correctly recovered, by comparing a second

check word derived from audio samples generated by said matrix multiplication subsystem against the check word extracted from the encoded bitstream.

14. An audio encoder configured to encode an N-channel audio program, wherein the program is specified over a time interval, the time interval includes a subinterval from a time t1 to a time t2, and a time-varying mix, A(t), of N encoded signal channels to M output channels has been specified over the time interval, where M is less than or equal to N, and wherein said encoder is configured to perform the method of any one of claims 1 to 4.

15. A decoder configured to implement recovery of an N-channel audio program, wherein the program is specified over a time interval, the time interval includes a subinterval from a time t1 to a time t2, and a time-varying mix, A(t), of N encoded signal channels to M output channels has been specified over the time interval, and wherein said decoder is configured to perform the method of any one of claims 5 to 13.