A Prototype Distributed Visualization System

Part III: Performance Benchmark

Contents

A Analysis of Benchmark Results **2**

B Details of Benchmark Results **10**

A Analysis of Benchmark Results

In the figures that follow, speed and compression ratio are plotted versus the number of nodes used in computation. For a description of how these results were obtained and definitions of the parameters and their options, please refer to Part I, Section 4, of this report.

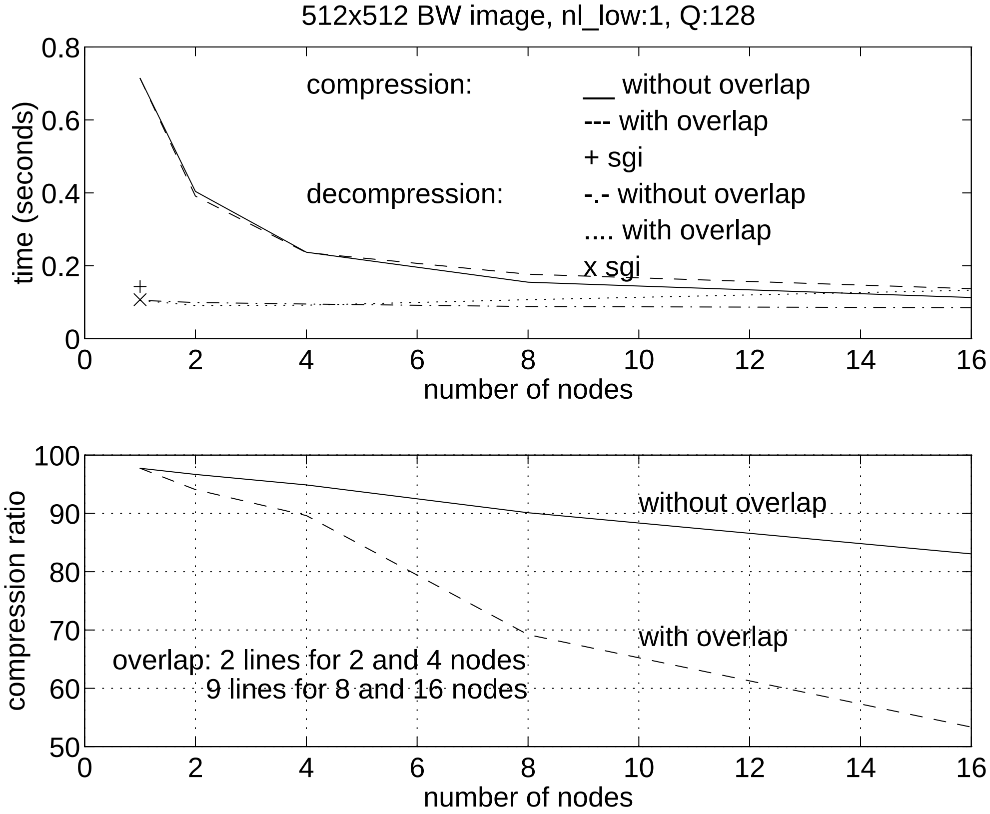

Figure 1: $nl_{low} = 1$, $Q = 2^7$

Figure 1 shows that the compression time on the Cray T3D decreases by less than 2 as the number of nodes increases by 2. For instance, 1 node achieves compression in 0.714 seconds, while 2 nodes do it in 0.404 seconds, which corresponds to a speedup of about 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, the speedup factor corresponding to entropy encoding is 1.1 at best. With 8 and 16 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. The compression speedup factor from 8 (0.155 seconds) to 16 nodes (0.113 seconds) is about 1.4. Compared to the compression time achieved with 1 node, the overall speedup factor provided by 16 nodes is 6.3.

Adding overlap to reduce discontinuity near inter-segment boundaries becomes a significant slowdown factor only when the amount is large relative to the size of the segment. This figure shows that the compression and decompression times remain practically unchanged when adding 2 lines of overlap at each inter-segment boundary, as is the case with 2 or 4 nodes. In contrast, there is a significant increase in compression and decompression times when using 8 and 16 nodes, since up to 9 lines of overlap are needed to reasonably approach the quality obtained with 1 node. Using 9 lines of overlap increases the compression time on 16 nodes by 21.2%.

Without overlap, the compression ratio (CR) obtained with 1 node is about 97.7 and about 83.1 with 16 nodes, which means a drop of 15.0% and a rate of decrease of about 1 per node. With 16 nodes and an overlap of 9 lines, the effective data area increases by 52.7% and the CR drops to 53.3, which is 45.4% less than the CR achieved by 1 node.

Using 16 nodes and overlap, compression is completed in 0.137 and decompression is achieved in 0.133 seconds. In comparison, the SGI achieves compression in 0.143 seconds and decompression in 0.107 seconds. Therefore, 16 nodes outperform the SGI on compression speed. However, since the SGI has an advantage on CR and image quality, the use of the Cray T3D is not justified in this case.

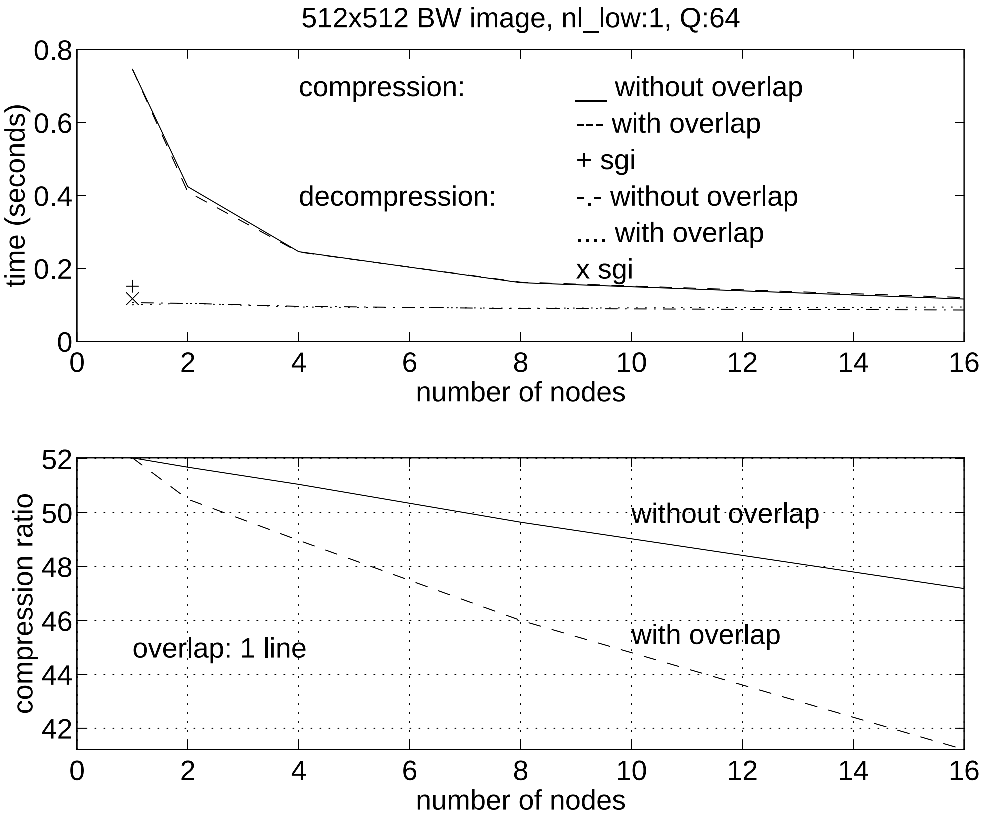

Figure 2: $nl_{low} = 1$, $Q = 2^6$

The reconstructed images corresponding to Figure 2 have better quality than those corresponding to the previous figure (Figure 1) because the quantization factor Q is lower. The speedup factors remain nearly the same. For instance, the compression speedup factor from 1 node (0.746 seconds) to 2 nodes (0.424 seconds) is 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, entropy encoding time decreases by a factor of about 1.1 as the number of nodes increases by 2. There is an increase in entropy encoding time with respect to the time shown in the previous figure since more bytes are required to encode less-quantized coefficients. With 8 and 16 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. Increasing the node count from 8 (0.161 seconds) to 16 (0.116 seconds) provides a compression speedup factor of only 1.4. Compared to the compression time achieved with 1 node, the overall speedup factor provided by 16 nodes is 6.4.

Without overlap, the CR obtained with 1 node is about 52.0 and 47.2 with 16 nodes, a drop of 9.3% and a rate of decrease of about 0.32 per node. With an overlap of 1 line at each inter-segment boundary, the CR drops from 50.5 (2 nodes) by about 0.66 per additional node used. So, the rate of decrease in CR has doubled by using 1 line of overlap at each boundary. With 16 nodes and this overlap, the effective data area increases by 5.9% and the CR drops to 41.2, which is 20.8% less than the CR achieved by 1 node.

Since the overlap amount used is small, there is practically no increase in compression or decompression time. Using 16 nodes and overlap, compression is completed in 0.120 and decompression is achieved in 0.094 seconds. In comparison, the SGI achieves compression in 0.151 seconds and decompression in 0.117 seconds. Therefore, 16 nodes outperform the SGI in terms of speed. In this case, the image quality obtained is superior to that obtained in the case of Figure 1 and the CR remains fairly high.

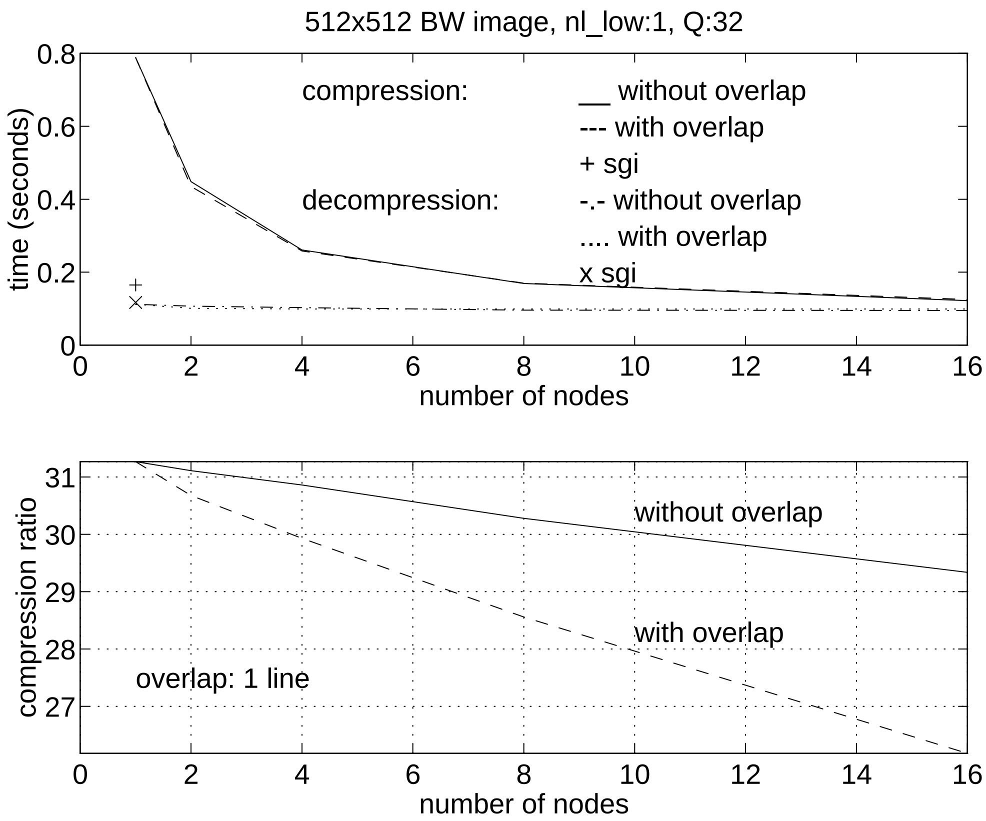

Figure 3: $nl_{low} = 1$, $Q = 2^5$

The reconstructed images corresponding to Figure 3 have better quality than those corresponding to the previous 2 figures (Figures 1 and 2) because the quantization factor Q is lower. The speedup factors remain nearly the same. For instance, 1 node achieves compression in 0.788 seconds, while 2 nodes do it in 0.448, which is equivalent to a speedup factor of 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, entropy encoding time decreases by a factor of about 1.2 as the number of nodes increases by 2. There is a slight increase in entropy encoding time with respect to the time shown in the previous figure since more bytes are required to encode less-quantized coefficients. With 8 and 16 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. Increasing the node count from 8 (0.169 seconds) to 16 (0.122 seconds) provides a compression speedup factor of only 1.4. Compared to the compression time achieved with 1 node, the overall speedup factor provided by 16 nodes is 6.5.

Without overlap, the CR obtained with 1 node is about 31.3 and 29.3 with 16 nodes, a drop of 6.2% and a rate of decrease of about 0.12 per node. With 1 line of overlap at each inter-segment boundary, the CR drops from 30.7 (2 nodes) at a rate of about 0.32 per additional node used. So, the rate of decrease in CR has more than doubled by using 1 line of overlap. With 16 nodes and this overlap, the effective data area increases by 5.9% and the CR drops to 26.2, which is 16.3% less than the CR achieved by 1 node.

Since the overlap amount used is small, there is practically no increase in compression or decompression time. Using 16 nodes and overlap, compression is completed in 0.125 and decompression is achieved in 0.099 seconds. In comparison, the SGI achieves compression in 0.165 seconds and decompression in 0.117 seconds. Therefore, 16 nodes outperform the SGI in terms of speed. In this case, the image quality obtained is better than in the cases of Figures 1 and 2. However, the CR may not be high enough if transmitting over a slow network.

In Figures 4, 5, and 6, the quantization factor is 256, 128, and 64, respectively. The compression parameters $nl_{low} = 0$ and $nl_{lh} = 4$. In terms of quality, the reconstructed images corresponding to Figures 4, 5, and 6 are comparable to the reconstructed images corresponding to Figures 1, 2, and 3, respectively. Again, as the quantization decreases, the quality of the reconstructed image increases.

Figure 4 shows that the compression time on the Cray T3D decreases by less than 2 as the number of nodes increases by 2. For instance, 1 node achieves compression in 1.071 seconds, while 2 nodes do it in 0.581 seconds. This is a speedup of about 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, the speedup factor corresponding to entropy encoding is 1.3 (from 1 to 2 and from 2 to 4 nodes) or 1.1 (from 4 to 8 and from 8 to 16 nodes). With 16 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. The compression speedup factor from 8 (0.200 seconds) to 16 nodes (0.133 seconds) is about 1.5. The speedup factor from 1 to 16 nodes is 8.1. Both of these factors are higher than those of Figures 1, 2, and 3. We can see that, unlike the curves in the earlier figures, Figure 4 shows that compression time falls below decompression time with 8 or more nodes. However, this is not so much due to the faster rate of decrease as it is due to the fact that decompression time is more than double what it was in these figures. The compression time has also increased with respect to these

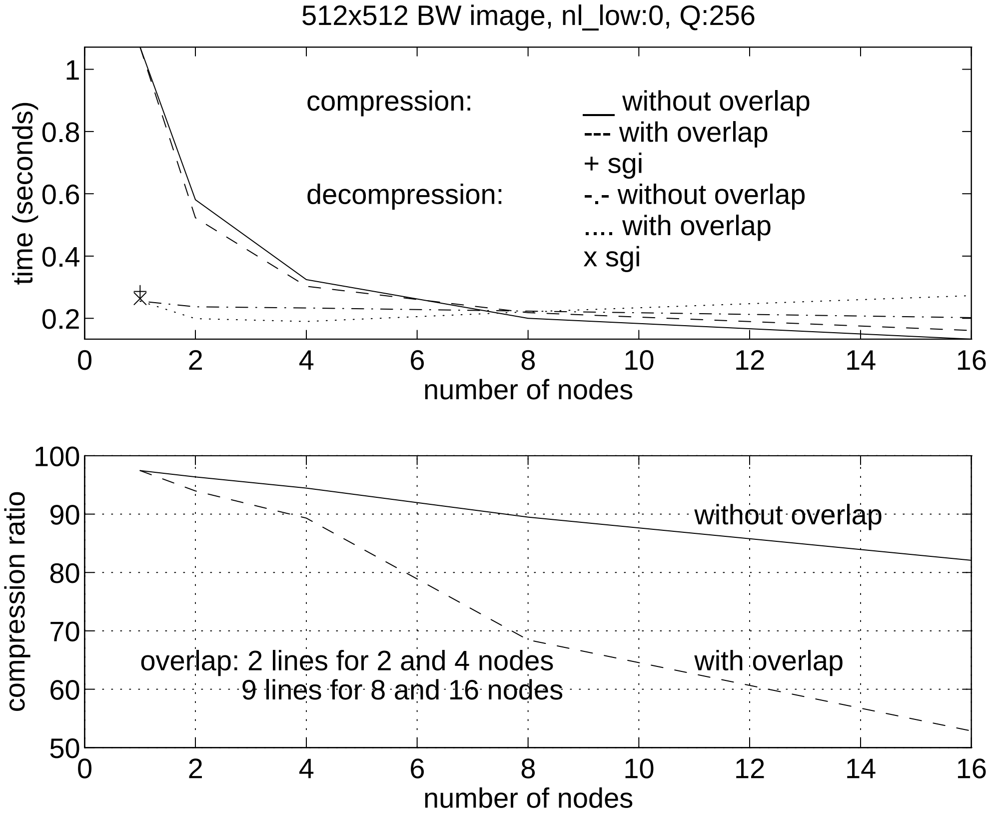

Figure 4: $nl_{low} = 0$, $Q = 2^8$

figures, but not by as much.

Adding overlap to reduce discontinuity near inter-segment boundaries becomes a significant slowdown factor only when the amount is large relative to the size of the segment. This figure shows that the compression and decompression times remain practically unchanged when adding 2 lines of overlap at each inter-segment boundary, as is the case with 2 or 4 nodes. In contrast, there is a significant increase in compression and decompression times when using 8 and 16 nodes, since up to 9 lines of overlap are needed to reasonably approach the quality obtained with 1 node. Using 9 lines of overlap increases the compression time on 16 nodes by 21.1%.

Without overlap, the compression ratio (CR) obtained with 1 node is about 97.5 and about 82.1 with 16 nodes, a drop of 15.8% and a rate of decrease of about 1 per node. With 16 nodes and an overlap of 9 lines, the effective data area increases by 52.7% and the CR drops to 52.9, which is 45.8% less than the CR achieved by 1 node.

Using 8 nodes and overlap, compression is completed in 0.218 and decompression is achieved in 0.221 seconds. In comparison, the SGI achieves compression in 0.286 seconds and decompression in 0.263 seconds. Therefore, 8 nodes outperform the SGI in terms of speed. However, since the SGI has an advantage on CR and image quality, the use of the Cray T3D may not be justified in this case.

The reconstructed images corresponding to Figure 5 have better quality than those cor-

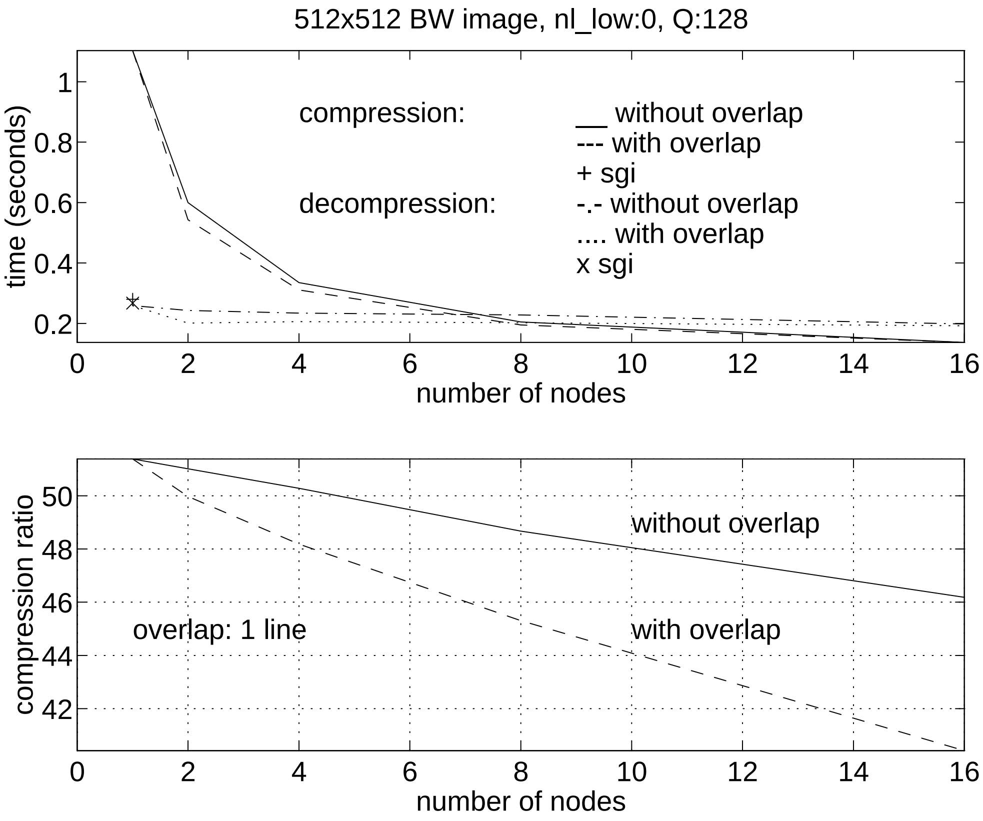

Figure 5: $nl_{low} = 0$, $Q = 2^7$

responding to the previous figure (Figure 4) because the quantization factor Q is lower. The speedup factors remain nearly the same. For instance, the compression speedup factor from 1 node (1.103 seconds) to 2 nodes (0.600 seconds) is 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, entropy encoding time decreases by a factor of about between 1.1 and 1.3 as the number of nodes increases by 2. There is an increase in entropy encoding time with respect to the time shown in the previous figure since more bytes are required to encode less-quantized coefficients. With 16 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. The compression speedup factor from 8 (0.205 seconds) to 16 nodes (0.137 seconds) is about 1.5. The speedup factor from 1 to 16 nodes is 8.1. Both of these factors are higher than those of Figures 1, 2, and 3. We can see that, unlike the curves in these figures, Figure 5 shows that compression time falls below decompression time with 8 or more nodes. However, this is not so much due to the faster rate of decrease as it is due to the fact that decompression time is more than double what it was in the earlier figures. The compression time has also increased, but not by as much.

The CR obtained with 1 node is 51.4. Without overlap, the CR obtained with 16 nodes is 46.2, a drop of 9.3% and a rate of decrease of about 0.32 per node. With an overlap of 1 line at each inter-segment boundary, the CR drops from 50.0 (2 nodes) by about 0.68 per node. So, the rate of decrease in CR has doubled by using 1 line of overlap at each boundary. With 16 nodes and this overlap, the effective data area increases by 5.9% and the CR drops to 40.4, which is 21.3% less than the CR achieved by 1 node.

Since the overlap amount added is small, there is practically no increase in compression or decompression time. Using 8 nodes and overlap, compression is completed in 0.195 and decompression is achieved in 0.202 seconds. In comparison, the SGI achieves compression in 0.280 seconds and decompression in 0.267 seconds. Therefore, 8 nodes outperform the SGI in terms of speed. In this case, the image quality obtained is superior to that obtained in the case of Figure 4 and the CR remains fairly high.

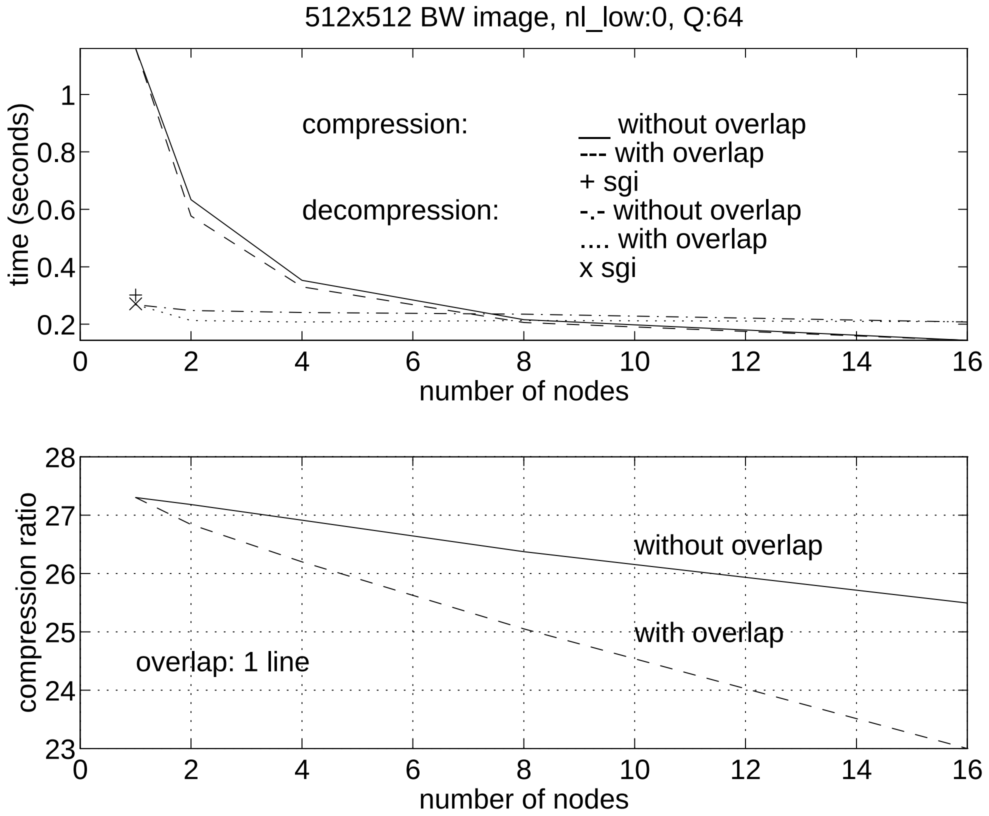

Figure 6: $nl_{low} = 0$, $Q = 2^6$

The reconstructed images corresponding to Figure 6 have better quality than those corresponding to the previous 2 figures (Figures 4 and 5) because the quantization factor Q is lower. The speedup factors remain nearly the same. For instance, 1 node achieves compression in 1.160 seconds, while 2 nodes do it in 0.634, which is equivalent to a speedup factor of 1.8. Both the FWT time and the quantization time decrease by a factor of 2 as the number of nodes increases by the same factor. In contrast, entropy encoding time decreases by a factor of about 1.4 or 1.2 as the number of nodes increases by 2. There is a slight increase in entropy encoding time with respect to the time shown in the previous figure since more bytes are required to encode less-quantized coefficients. With 8 nodes, entropy encoding time is greater than FWT and quantization times combined. Therefore, using more nodes becomes less and less efficient. The compression speedup factor from 8 (0.216 seconds) to 16 nodes (0.144 seconds) is about 1.5. The speedup factor from 1 to 16 nodes is 8.1. Both of these factors are higher than those of Figures 1, 2, and 3. We can see that, unlike the curves in these figures, Figure 5 shows that compression time falls below decompression time with 8 or more nodes. However, this is not so much due to the faster rate of decrease as it is due to the fact that decompression time is more than double what it was in these figures. The compression time has also increased, but not by as much.

The CR obtained with 1 node is 27.3. Without overlap, the CR obtained with 16 nodes is 25.5, which means a drop of 6.6% and a rate of decrease of about 0.11 per node. With 1 line of overlap at each inter-segment boundary, the CR drops from 26.8 (2 nodes) at a rate of about 0.27 per additional node used. So, the rate of decrease in CR has doubled by using 1 line of overlap. With 16 nodes and this overlap, the effective data area increases by 5.9% and the CR drops to 23.0, which is 15.8% less than the CR achieved by 1 node.

Since the overlap amount used is small, there is practically no increase in compression or decompression time. Using 8 nodes and overlap, compression is completed in 0.206 and decompression is achieved in 0.213 seconds. In comparison, the SGI achieves compression in 0.302 seconds and decompression in 0.271 seconds. Therefore, 8 nodes outperform the SGI in terms of speed. In this case, the image quality obtained is better than in the cases of Figures 4 and 5. However, the CR may not be high enough if transmitting over a slow network.

B Details of Benchmark Results

The tables below provide details of the benchmark results, as described in Part I, Section 4, of this report. In these tables, the timing data for the compression steps represent the average value over ten iterations, followed by the minimum value (in parenthesis). The data for the decompression steps represent the minimum value over ten iterations.

Table 1: without overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^7$

Number of PEs	1	2	4	8	16
Compress	0.714	0.404 (0.402)	0.237 (0.235)	0.155 (0.154)	0.113 (0.112)
FWT	0.557	0.281 (0.278)	0.137 (0.136)	0.069 (0.068)	0.033
Quantize	0.032	0.016	0.008	0.004	0.002
Entropy encode	0.098	0.095	0.086 (0.084)	0.080 (0.078)	0.077 (0.075)
RLEncode	0.033	0.018 (0.015)	0.009 (0.007)	0.005 (0.003)	0.003 (0.002)
Huffman	0.065	0.079 (0.077)	0.077 (0.076)	0.075 (0.074)	0.074
Histogram gen	0.0037	0.0025 (0.0017)	0.0013 (0.0007)	0.0008 (0.0002)	0.0004 (0.0001)
Ghistogram gen	0.0000	0.014 (0.008)	0.017 (0.014)	0.017 (0.015)	0.017 (0.016)
HTable gen	0.025	0.030	0.030	0.030	0.030
HEncode	0.016	0.010 (0.007)	0.005 (0.003)	0.003 (0.001)	0.001
Decompress	0.105	0.099	0.095	0.088	0.085
IFWT	0.081	0.076	0.072	0.064	0.060
IQuantize	0.009	0.009	0.009	0.009	0.006
Entropy decode	0.013	0.012	0.012	0.011	0.010
RLE data size	4992	5005	5037	5171	5331
HTable size	131	131	127	127	123
HE data size	2527	2536	2552	2618	2709
CR	97.742	96.696	94.877	90.115	83.062
RMS error	11.836	11.960	12.073	12.267	12.789

Table 2: without overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^6$

Number of PEs	1	2	4	8	16
Compress	0.746	0.424 (0.418)	0.246 (0.241)	0.161 (0.158)	0.116 (0.114)
FWT	0.558	0.281 (0.280)	0.137 (0.136)	0.069 (0.068)	0.033
Quantize	0.032	0.016	0.008	0.004	0.002
Entropy encode	0.130	0.115 (0.108)	0.095 (0.090)	0.086 (0.082)	0.080 (0.078)
RLEncode	0.044	0.024 (0.020)	0.013 (0.008)	0.007 (0.004)	0.004 (0.002)
Huffman	0.086	0.090 (0.088)	0.083 (0.082)	0.079 (0.078)	0.077 (0.076)
Histogram gen	0.0081	0.0051 (0.0034)	0.0027 (0.0012)	0.0015 (0.0004)	0.0008 (0.0002)
Ghistogram gen	0.0000	0.015 (0.009)	0.019 (0.013)	0.019 (0.015)	0.018 (0.016)
HTable gen	0.026	0.031 (0.030)	0.031 (0.030)	0.031 (0.030)	0.031 (0.030)
HEncode	0.032	0.019 (0.013)	0.010 (0.005)	0.005 (0.002)	0.003 (0.001)
Decompress	0.106	0.104	0.096	0.090	0.086
IFWT	0.078	0.078	0.070	0.063	0.058
IQuantize	0.009	0.009	0.008	0.008	0.007
Entropy decode	0.017	0.016	0.016	0.016	0.016
RLE data size	9610	9647	9694	9804	10045
HTable size	203	203	199	203	203
HE data size	4811	4825	4852	4913	5029
CR	52.033	51.685	51.050	49.648	47.182
RMS error	8.683	8.731	8.789	8.891	9.111

Table 3: without overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^5$

Number of PEs	1	2	4	8	16
Compress	0.788	0.448 (0.439)	0.261 (0.254)	0.169 (0.164)	0.122 (0.118)
FWT	0.558	0.279 (0.278)	0.137	0.069 (0.068)	0.033
Quantize	0.032	0.016	0.008	0.004	0.002
Entropy encode	0.173	0.139 (0.132)	0.109 (0.102)	0.093 (0.089)	0.085 (0.082)
RLEncode	0.059	0.032 (0.026)	0.017 (0.011)	0.009 (0.005)	0.005 (0.002)
Huffman	0.114	0.106 (0.105)	0.093 (0.091)	0.085 (0.084)	0.081 (0.080)
Histogram gen	0.0132	0.0085 (0.0061)	0.0045 (0.0022)	0.0024 (0.0008)	0.0013 (0.0003)
Ghistogram gen	0.0000	0.018 (0.008)	0.021 (0.013)	0.021 (0.015)	0.019 (0.016)
HTable gen	0.027	0.032	0.032	0.032	0.032
HEncode	0.053	0.031 (0.023)	0.016 (0.009)	0.008 (0.003)	0.005 (0.001)
Decompress	0.112	0.107	0.103	0.096	0.095
IFWT	0.080	0.076	0.069	0.062	0.059
IQuantize	0.009	0.009	0.008	0.008	0.006
Entropy decode	0.022	0.021	0.021	0.022	0.021
RLE data size	16028	16071	16129	16264	16478
HTable size	347	347	343	347	339
HE data size	8013	8035	8068	8147	8273
CR	31.267	31.111	30.859	30.278	29.336
RMS error	6.622	6.646	6.682	6.759	6.888

Table 4: without overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^8$

Number of PEs	1	2	4	8	16
Compress	1.071	0.581 (0.579)	0.324 (0.322)	0.200 (0.198)	0.133 (0.132)
FWT	0.672	0.336	0.165	0.082	0.038
Quantize	0.128	0.064	0.032	0.016	0.008
Entropy encode	0.167	0.129 (0.127)	0.101 (0.099)	0.088 (0.086)	0.081 (0.080)
RLEncode	0.102	0.052 (0.050)	0.026 (0.024)	0.013 (0.012)	0.007 (0.006)
Huffman	0.065	0.077	0.075	0.075 (0.074)	0.074
Histogram gen	0.0037	0.0025 (0.0017)	0.0013 (0.0007)	0.0008 (0.0002)	0.0004 (0.0001)
Ghistogram gen	0.000	0.012 (0.008)	0.015 (0.012)	0.016 (0.014)	0.017 (0.016)
HTable gen	0.025	0.030	0.030	0.030	0.030
HEncode	0.017	0.010 (0.007)	0.005 (0.003)	0.003 (0.001)	0.002 (0.001)
Decompress	0.255	0.237	0.233	0.223	0.202
IFWT	0.171	0.160	0.154	0.146	0.111
IQuantize	0.042	0.038	0.038	0.034	0.041
Entropy decode	0.041	0.037	0.036	0.037	0.037
RLE data size	4982	4992	5022	5163	5348
HTable size	139	139	135	135	131
HE data size	2527	2537	2556	2630	2738
CR	97.451	96.376	94.466	89.499	82.100
RMS error	11.827	11.951	12.070	12.271	12.750

Table 5: without overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^7$

Number of PEs	1	2	4	8	16
Compress	1.103	0.600 (0.595)	0.335 (0.330)	0.205 (0.202)	0.137 (0.135)
FWT	0.672	0.336	0.165	0.082	0.038
Quantize	0.128	0.064	0.032	0.016	0.008
Entropy encode	0.199	0.148 (0.143)	0.112 (0.107)	0.094 (0.090)	0.084 (0.082)
RLEncode	0.113	0.059 (0.054)	0.030 (0.026)	0.015 (0.012)	0.008 (0.006)
Huffman	0.086	0.089 (0.088)	0.082 (0.081)	0.079 (0.078)	0.077 (0.076)
Histogram gen	0.0081	0.0052 (0.0034)	0.0027 (0.0012)	0.0015 (0.0004)	0.0008 (0.0002)
Ghistogram gen	0.0000	0.015 (0.008)	0.018 (0.012)	0.018 (0.014)	0.018 (0.016)
HTable gen	0.026	0.031	0.031	0.031	0.031
HEncode	0.032	0.019 (0.013)	0.010 (0.005)	0.005 (0.002)	0.003 (0.001)
Decompress	0.259	0.243	0.234	0.228	0.198
IFWT	0.172	0.162	0.151	0.146	0.115
IQuantize	0.040	0.038	0.037	0.036	0.037
Entropy decode	0.044	0.041	0.041	0.040	0.038
RLE data size	9674	9713	9774	9901	10161
HTable size	219	219	215	215	211
HE data size	4859	4876	4915	5007	5141
CR	51.381	51.011	50.277	48.671	46.185
RMS error	8.652	8.702	8.771	8.877	9.108

Table 6: without overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^6$

Number of PEs	1	2	4	8	16
Compress	1.160	0.634 (0.624)	0.353 (0.344)	0.216 (0.209)	0.144 (0.140)
FWT	0.672	0.336	0.165	0.082	0.038
Quantize	0.128	0.064	0.032	0.016	0.008
Entropy encode	0.256	0.182 (0.172)	0.131 (0.121)	0.105 (0.098)	0.091 (0.087)
RLEncode	0.131	0.070 (0.062)	0.036 (0.028)	0.018 (0.013)	0.010 (0.007)
Huffman	0.124	0.112 (0.110)	0.095 (0.093)	0.086 (0.085)	0.082 (0.0081)
Histogram gen	0.0154	0.0096 (0.0067)	0.0051 (0.0023)	0.0028 (0.0008)	0.0015 (0.0003)
Ghistogram gen	0.0000	0.019 (0.008)	0.023 (0.012)	0.022 (0.014)	0.020 (0.016)
HTable gen	0.027	0.032	0.032	0.032	0.032
HEncode	0.061	0.036 (0.025)	0.019 (0.009)	0.010 (0.003)	0.005 (0.001)
Decompress	0.268	0.248	0.241	0.235	0.208
IFWT	0.171	0.161	0.149	0.145	0.110
IQuantize	0.041	0.038	0.038	0.038	0.045
Entropy decode	0.052	0.049	0.049	0.048	0.047
RLE data size	18190	18243	18337	18530	18855
HTable size	367	367	363	367	359
HE data size	9210	9233	9293	9408	9600
CR	27.304	27.182	26.914	26.375	25.493
RMS error	6.312	6.332	6.365	6.432	6.563

Table 7: with overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^7$

Number of PEs	1	2, overlap:2	4, overlap:2	8, overlap:9	16, overlap:9
Compress	0.714	0.391 (0.388)	0.236 (0.234)	0.177 (0.175)	0.137 (0.135)
FWT	0.557	0.267 (0.266)	0.136 (0.134)	0.086 (0.076)	0.051 (0.044)
Quantize	0.032	0.016	0.008	0.005	0.003
Entropy encode	0.098	0.095 (0.092)	0.086 (0.084)	0.093 (0.080)	0.088 (0.078)
RLEncode	0.033	0.018 (0.016)	0.009 (0.007)	0.006 (0.004)	0.004 (0.002)
Huffman	0.065	0.077 (0.076)	0.077 (0.075)	0.088 (0.075)	0.085 (0.075)
Histogram gen	0.0037	0.0026 (0.0018)	0.0014 (0.0007)	0.0010 (0.0003)	0.0006 (0.0002)
Histogram gen	0.0000	0.011 (0.008)	0.017 (0.012)	0.029 (0.014)	0.028 (0.016)
HTable gen	0.025	0.030	0.030	0.030	0.030
HEncode	0.016	0.010 (0.007)	0.005 (0.003)	0.004 (0.001)	0.002 (0.001)
Decompress	0.105	0.091	0.092	0.107	0.133
IFWT	0.081	0.068	0.067	0.079	0.094
IQuantize	0.009	0.009	0.008	0.011	0.016
Entropy decode	0.013	0.013	0.013	0.015	0.020
RLE data size	4992	5122	5306	6791	8591
HTable size	131	135	131	135	135
HE data size	2527	2607	2710	3490	4456
CR	97.742	94.093	89.622	69.186	53.336
RMS error	11.836	11.726	11.713	11.698	11.660

Table 8: with overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^6$

Number of PEs	1	2, overlap:1	4, overlap:1	8, overlap:1	16, overlap:1
Compress	0.746	0.410 (0.405)	0.245 (0.240)	0.162 (0.158)	0.120 (0.118)
FWT	0.558	0.267	0.134	0.068 (0.067)	0.035 (0.034)
Quantize	0.032	0.016	0.008	0.004	0.002
Entropy encode	0.130	0.114 (0.109)	0.096 (0.091)	0.086 (0.083)	0.081 (0.079)
RLEncode	0.044	0.025 (0.020)	0.013 (0.009)	0.007 (0.004)	0.004 (0.002)
Huffman	0.086	0.090 (0.089)	0.083 (0.082)	0.080 (0.079)	0.078 (0.077)
Histogram gen	0.0081	0.0052 (0.0035)	0.0028 (0.0013)	0.0016 (0.0005)	0.0009 (0.0002)
Ghistogram gen	0.0000	0.015 (0.008)	0.019 (0.012)	0.019 (0.014)	0.019 (0.016)
HTable gen	0.026	0.031	0.031 (0.030)	0.031 (0.030)	0.031
HEncode	0.032	0.019 (0.013)	0.010 (0.005)	0.006 (0.002)	0.003 (0.001)
Decompress	0.106	0.100	0.094	0.091	0.094
IFWT	0.078	0.071	0.066	0.064	0.064
IQuantize	0.009	0.010	0.009	0.008	0.008
Entropy decode	0.017	0.016	0.016	0.017	0.017
RLE data size	9610	9803	9979	10400	11193
HTable size	203	207	207	207	211
HE data size	4811	4940	5062	5327	5826
CR	52.033	50.500	48.971	46.006	41.211
RMS error	8.683	8.580	8.545	8.482	8.386

Table 9: with overlap, $nl_{low} = 1$, $nl_{lh} = 3$, $Q = 2^5$

Number of PEs	1	2, overlap:1	4, overlap:1	8, overlap:1	16, overlap:1
Compress	0.788	0.435 (0.428)	0.258 (0.251)	0.170 (0.165)	0.125 (0.122)
FWT	0.558	0.267	0.134 (0.133)	0.068 (0.067)	0.036 (0.034)
Quantize	0.032	0.016	0.008	0.004	0.002
Entropy encode	0.173	0.139 (0.131)	0.109 (0.103)	0.094 (0.090)	0.086 (0.083)
RLEncode	0.059	0.032 (0.027)	0.017 (0.011)	0.009 (0.005)	0.005 (0.002)
Huffman	0.114	0.106 (0.105)	0.093 (0.091)	0.086 (0.085)	0.082 (0.080)
Histogram gen	0.0132	0.0084 (0.0062)	0.0045 (0.0023)	0.0025 (0.0008)	0.0014 (0.0004)
Ghistogram gen	0.0000	0.016 (0.008)	0.021 (0.012)	0.022 (0.015)	0.020 (0.016)
HTable gen	0.027	0.032	0.032	0.032	0.032
HEncode	0.053	0.031 (0.023)	0.016 (0.009)	0.009 (0.004)	0.005 (0.002)
Decompress	0.112	0.101	0.099	0.099	0.099
IFWT	0.080	0.069	0.067	0.064	0.063
IQuantize	0.009	0.010	0.009	0.008	0.007
Entropy decode	0.022	0.021	0.021	0.022	0.022
RLE size	16028	16157	16416	16871	17891
HTable size	347	351	347	351	355
HE data size	8013	8150	8328	8664	9334
CR	31.267	30.678	29.929	28.559	26.180
RMS error	6.622	6.589	6.583	6.555	6.511

Table 10: with overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^8$

Number of PEs	1	2, overlap:2	4, overlap:2	8, overlap:9	16, overlap:9
Compress	1.071	0.523 (0.520)	0.303 (0.301)	0.218 (0.217)	0.161 (0.160)
FWT	0.672	0.276	0.140 (0.138)	0.088 (0.077)	0.053 (0.045)
Quantize	0.128	0.064	0.033	0.021 (0.018)	0.012 (0.010)
Entropy encode	0.167	0.130 (0.127)	0.105 (0.102)	0.108 (0.091)	0.097 (0.084)
RLEncode	0.102	0.052 (0.050)	0.027 (0.025)	0.017 (0.014)	0.011 (0.008)
Huffman	0.065	0.077	0.079 (0.075)	0.093 (0.075)	0.089 (0.075)
Histogram gen	0.0037	0.0026 (0.0018)	0.0014 (0.0007)	0.0010 (0.0003)	0.0006 (0.0002)
Ghistogram gen	0.000	0.011 (0.008)	0.018 (0.012)	0.035 (0.014)	0.031 (0.016)
HTable gen	0.025	0.030	0.030	0.030	0.030
HEncode	0.017	0.010 (0.007)	0.005 (0.003)	0.004 (0.001)	0.002 (0.001)
Decompress	0.255	0.199	0.190	0.221	0.273
IFWT	0.171	0.119	0.111	0.123	0.150
IQuantize	0.042	0.039	0.038	0.047	0.055
Entropy decode	0.041	0.038	0.036	0.046	0.056
RLE data size	4982	5102	5286	6828	8632
HTable size	139	143	139	147	139
HE data size	2527	2603	2712	3519	4496
CR	97.451	93.958	89.317	68.445	52.862
RMS error	11.827	11.711	11.705	11.713	11.659

Table 11: with overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^7$

Number of PEs	1	2, overlap:1	4, overlap:1	8, overlap:1	16, overlap:1
Compress	1.103	0.543 (0.537)	0.311 (0.306)	0.195 (0.191)	0.137 (0.135)
FWT	0.672	0.276	0.139 (0.138)	0.070 (0.069)	0.036 (0.035)
Quantize	0.128	0.064	0.033 (0.032)	0.017 (0.016)	0.009 (0.008)
Entropy encode	0.199	0.150 (0.143)	0.113 (0.109)	0.095 (0.092)	0.086 (0.084)
RLEncode	0.113	0.059 (0.054)	0.030 (0.026)	0.016 (0.013)	0.009 (0.006)
Huffman	0.086	0.091 (0.090)	0.084 (0.083)	0.080 (0.079)	0.079 (0.077)
Histogram gen	0.0081	0.0052 (0.0035)	0.0028 (0.0013)	0.0016 (0.0005)	0.0009 (0.0002)
Ghistogram gen	0.0000	0.015 (0.008)	0.019 (0.012)	0.020 (0.014)	0.020 (0.016)
HTable gen	0.026	0.031	0.031	0.031	0.031
HEncode	0.032	0.020 (0.014)	0.010 (0.005)	0.006 (0.002)	0.003 (0.001)
Decompress	0.259	0.201	0.206	0.202	0.192
IFWT	0.172	0.119	0.119	0.117	0.103
IQuantize	0.040	0.038	0.040	0.039	0.040
Entropy decode	0.044	0.042	0.043	0.043	0.045
RLE data size	9674	9845	10067	10487	11333
HTable size	219	223	223	219	219
HE data size	4859	4979	5134	5403	5942
CR	51.381	49.970	48.179	45.307	40.423
RMS error	8.652	8.568	8.519	8.457	8.366

Table 12: with overlap, $nl_{low} = 0$, $nl_{lh} = 4$, $Q = 2^6$

Number of PEs	1	2, overlap:1	4, overlap:1	8, overlap:1	16, overlap:1
Compress	1.160	0.577 (0.566)	0.330 (0.320)	0.206 (0.199)	0.144 (0.140)
FWT	0.672	0.276	0.139 (0.138)	0.070 (0.069)	0.036 (0.035)
Quantize	0.128	0.064	0.033 (0.032)	0.017 (0.016)	0.009 (0.008)
Entropy encode	0.256	0.183 (0.173)	0.132 (0.123)	0.106 (0.099)	0.092 (0.089)
RLEncode	0.131	0.070 (0.062)	0.036 (0.029)	0.019 (0.014)	0.010 (0.007)
Huffman	0.124	0.114 (0.110)	0.096 (0.094)	0.087 (0.086)	0.083 (0.0081)
Histogram gen	0.0154	0.0096 (0.0068)	0.0052 (0.0023)	0.0029 (0.0008)	0.0016 (0.0004)
Ghistogram gen	0.0000	0.018 (0.008)	0.023 (0.012)	0.023 (0.014)	0.021 (0.015)
HTable gen	0.027	0.032	0.032	0.032	0.033 (0.032)
HEncode	0.061	0.037 (0.026)	0.019 (0.009)	0.011 (0.003)	0.006 (0.002)
Decompress	0.268	0.213	0.208	0.213	0.209
IFWT	0.171	0.123	0.116	0.120	0.105
IQuantize	0.041	0.038	0.040	0.039	0.047
Entropy decode	0.052	0.049	0.050	0.050	0.051
RLE data size	18190	18337	18617	19161	20326
HTable size	367	371	371	375	375
HE data size	9210	9352	9550	9925	10699
CR	27.304	26.840	26.201	25.052	22.999
RMS error	6.312	6.290	6.277	6.244	6.195