Telairity Dives Deep into 4K Technology – Part 11

Following the advice of the 1968 Jerry Lewis film, "Don't Raise the Bridge, Lower the River", MPEG-5 compression presumably offers a way to lessen the flood of bits created by shifting from 2K HD to 4K UHD resolution. Moreover, looking at the 10-year pattern followed by previous releases of new video compression standards, it might be reasonable to expect still more advanced MPEG-6 compression around 2023, presumably just in time to handle the next 10X increase in bits, created by shifting from 4K UHD to 8K UHD resolution.

Regrettably for this optimistic scenario, new compression standards are forced to follow a path of diminishing gains. The logic here is not hard to see. Unlike Moore's Law (mentioned earlier in this series as the prototype for seeming endless improvements in capacity), compression counts down rather than up. Counting up—in the case of Moore's Law, doubling transistor counts every couple of years—pushes against an elastic ceiling with no obvious limit to its expansion capabilities.

Alas, the logic of counting down is very different. Rather than an indefinitely elastic ceiling above, there is a hard floor below. Recall the two-fold strategy that underlies digital compression. First, eliminate everything redundant ("lossless" compression). Second, when everything redundant is gone, eliminate everything unimportant ("lossy" compression). When all the unimportant data (however defined) is gone, the limit of what can be done—without comprising quality—has been reached.

Moreover, the closer the approach to the fixed floor set by the pool of all and only significant data, the higher the resistance to further gains. Using relatively modest compute power, MPEG-2 takes 1500 Mbps of raw HD video down to 15 Mbps—a reduction of 99%. By throwing a lot more computational resources at the task of compression, MPEG-4 eliminates a further two-thirds of the remaining data, leaving a residue of just 5 Mbps.

By throwing an extravagant amount of computational energy at the task, MPEG-5 *might* be able to eliminate half of what remains, reducing 5 Mbps to 2.5 Mbps. (At least, that is the goal set for MPEG-5, though current demonstrations of the technology generally settle for reductions of 20-30% rather than the postulated 50%.) MPEG-6, should it be developed, would have to push far harder still to eke out any additional gains, possibly going below 2 Mbps. And it is likely that no amount whatsoever of computational power could ever push HD video below 1.5 Mbps (setting the hard floor at 1 significant bit in every 1000).

Note the sharply diminishing returns from additional compression. The great majority of what can be achieved by compression is achieved by MPEG-2, eliminating 99 out of every 100 bits. MPEG-4 manages a substantially additional gain by tossing out 2 of every 3 remaining bits. Realistically speaking, at this point, the bit stream has largely been wrung dry. Only 1 bit in 300 remains, and all of these cannot possibly be eliminated. Worse, in absolute terms, even if the 5 remaining Mbps went to zero, the savings would still be only half the 10 Mbps reduction realized in the previous step. The following figure illustrates this curve of sharply diminishing results for dramatically escalating efforts.

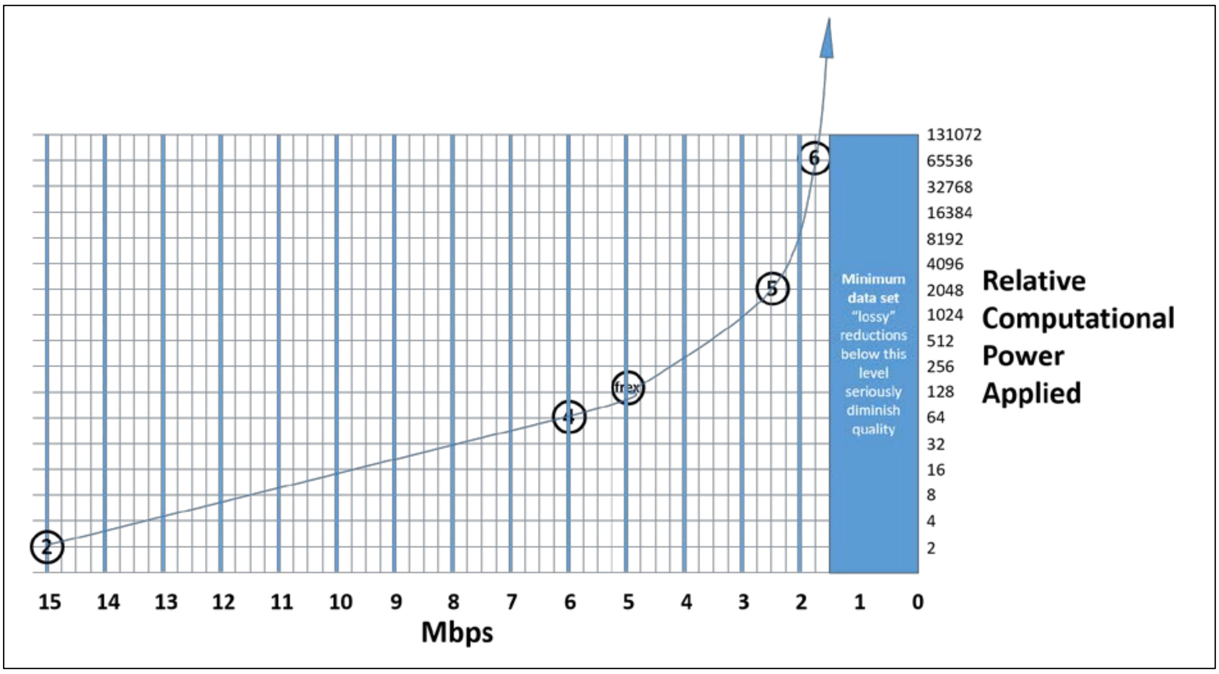

The compression curve for HD video, from MPEG-2 to a hypothetical MPEG-6. As compression approaches the floor set by the amount of non-redundant, non-trivial data (generously, if arbitrarily, fixed here at 1 bit in 1000), real gains decrease while the amount of effort required to make further progress rises exponentially. The assumption used in this graph is that successive MPEG standards appear at 10 year intervals, allowing time for 5 "Moore's Law" doublings in available computational power. This means each new standard has roughly 32X the power of the previous standard to throw at the compression problem, starting from an arbitrary level of 2 for MPEG-2.

To be sure, the above numbers are merely illustrative, and lack both scientific authority and general validity. But the basic point they make is inescapable. There is hard floor to what can be achieved with data compression, and each successive step taken in the direction of that floor will achieve less in the way of results while requiring more in the way of effort. In truth, the last really dramatic step possible with compression technology was MPEG-4. MPEG-5, and any further video compression standards that may be developed, will not be dramatic leaps forward from MPEG-4, but rather increasingly modest gains.

Given this logic of diminishing returns, what is the bottom line for 4K UHD video? As we have seen, the increase in bits required by new UHD resolution standards is a compounding problem, composed of three independent terms: more frames per second, more pixels per frame, and more bits per pixel. When fully realized, this equation multiplies out to an order of magnitude increase for 4K UHD over 2K HD, raising raw data rates from 1.5 to 15 Gbps.

Now apply the above rules for successive compression levels to a 15 Gbps UHD data stream. Eliminating 99 out of every 100 bits with MPEG-2 technology leaves a staggering 150 Mbps. Eliminating 2 out of every 3 remaining bits with MPEG-4 reduces this to a still hefty 50 Mbps. MPEG-5—depending on how close it gets to the target of eliminating 1 of every 2 bits left by MPEG-4—will be able to further thin the UHD stream down to 25-35 Mbps.

How bad is this? We will pursue that answer in the next part of this series.