Comment to Energy Star – STB

Testing procedure and power distribution

Background

- Energy Star-STB defines base and additional power allowances using average STB power measurements;
- The testing procedure defined by Energy Star-STB is random in essence: a very small number of boxes are randomly picked and tested. This inherently assumes that any STB is representative of all the STBs within a line of products. As a consequence, the worst STB (power wise) has to meet Energy Star requirements that were defined based on average numbers.
  In our view, there's a contradiction.

Chip power & distribution

The short description below applies to chips (in general) but may also apply to other types of parts (hard drives, supplies…) present in an STB.

Any 2 parts of the very same chip design (same part number) have differences: due to the fabrication process, there are variations down the transistor level.
One consequence is that 2 parts of the same chip design, can have different power characteristics. A distribution of the parts can be computed: most of the parts will be "typical" and burn an average power, some will be "slower" and burn less power, and, finally, some will be "faster" and burn more power.

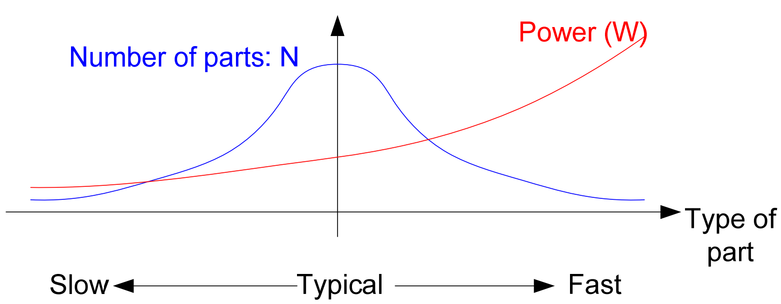

As shown in the diagram above, most parts tend to be typical. The power on the other hand tends to increase "exponentially" for fast parts. A very small number of fast parts have power characteristics that are noticeably higher than typical parts.

Parts are screened at production, but the [slow-typical-fast] window remains large: the parts that are kept work well. Rejecting more parts would be wasting perfectly valid parts and thus increasing the cost.

The increase in power in fast parts can be measured while a chip is on (impact on P_{TV} of the Energy Star STB specification) or idle (P_{SLEEP} on the Energy Star STB specification).

Power consumed in a chip has multiple components. One of them is generally referred to as "leakage". Leakage is present in a chip whenever it is powered. It increases dramatically for fast parts.
To give a ballpark estimate, an STB with a fast main SoC (processor/video/audio…) may burn about 1W (AC) more than an STB with the same, but typical, part, while the STB is in standby.
As technology shrinks (65nm, 45nm, 32nm…), the share of leakage in the total power increases.

<u>Conclusion</u>

- The point here isn't to discuss solutions that a chip/STB maker can develop to alleviate the power distribution issue. There are some. In general they add to the cost (likely passed to the consumer) of the final product and can only decrease partially the power spread.
- As explained above, the testing procedure imposes that all STBs within a product line fit within the power budget: an STB using a fast part has a very low probability to exist and be picked for testing. Yet if it is picked and tested, its higher power may exceed the Energy Star targets (computed with average cases in mind) and cause the entire product line to fail, even though say more than 99% of the STBs would have passed.
- The goal of the following suggestions is to make the testing procedure less sensitive to rare "outliers".
    - Pick STBs from different lots (or with parts from different lots);
    - Allow the tester to reject one or more outlying (high) measurement, with, possibly, an increase of the number of tested STBs by one or more.
    - Alternatively, instead of rejecting outright the outlying measurement(s), it/they could be allowed to exceed a kWh/y target by a fixed margin. Other measurements would be required to be below the target.