Workload Frequency Scaling law—

Derivation and Verification.

NOOR MUBEEN, INTEL

**WORKLOAD
SCALABILITY
HAS A CASCADE
RELATION VIA THE
SCALE FACTOR.**

Many processors expose performance-monitoring counters that help measure 'productive performance' associated with workloads. Productive performance is typically represented by *scale factor*, a term that refers to the extent of *stalls* compared with *stall-free* cycles within a time window. The scale factor of workload is also influenced by clock frequency as selected by frequency-selection governors. Hence, in a dynamic voltage/frequency scaling or DVFS system (such as Intel Speed Shift¹), the utilization, power, and performance outputs are also functions of the scale factor and its variations. Some

governance algorithms do treat the scale factor in ways that are native to their governance philosophy.

This article presents equations that relate to workload utilization scaling at a per-DVFS subsystem level. A relation between frequency, utilization, and scale factor (which itself varies with frequency) is established. The verification of these equations turns out to be tricky, since inherent to workload, the utilization also varies seemingly in an unspecified manner at the granularity of governance samples. Thus, a novel approach called *histogram ridge trace* is applied. Quantifying the scaling impact is critical when treating DVFS as a building block. Typical application includes DVFS governors and other layers that influence utilization, power, and performance of the system. The scope here though, is limited to demonstrating well-quantified and verified scaling equations.

WORKLOAD SCALING

Intel has three architecture-independent registers that are relevant to this topic:

- A_{perf} A running counter that counts at actual clock rate of execution at that instant. This actual clock frequency may vary over time based on governance and/or other algorithms. This register counts only during active state (CO).
- M_{perf} A running counter for activity that counts at a fixed TSC (time stamp counter) clock rate. It counts only during active state (CO).
- P_{perf} This counter is similar to A_{perf} except that it does not count when the activity stalls as a result of some dependency, likely gated on another IP's clock domain (e.g.,

memory].

The delta of these counters in a given time window commonly interpret the following:

- Utilization. $U = \Delta M_{perf} / \Delta TSC$
- Scale factor. $S = \Delta P_{perf} / \Delta A_{perf}$

Ideally, if an activity is free of stalls, then $\Delta P_{perf} = \Delta A_{perf}$ (i.e., the scale factor S will be 1). In such a situation, the time taken to complete an activity in a given time window would simply be the inverse of the actual frequency in that window. In most real workloads the scale factor varies and is often less than 1. Establishing accurate equations relating utilization, scale factor, and frequency allows DVFS governors to dispense well-informed frequency-change decisions.

At the outset, it may seem that the workload-scaling problem is similar to the well-known Amdahl's law, which deals with workload speedup as associated with parallel-compute resources. However, Amdahl's law cannot be applied here since the “parallelizability factor” analogy to “scale factor” doesn't hold up—the former is independent of the resource being scaled (i.e., parallel processors), while the latter is a function of the resource being scaled (i.e., frequency). Utilization scalability over frequency has a cascade relation via the scale factor.

TERMINOLOGY AND ASSUMPTIONS

This article makes use of following terminology:

- *Utilization or CO activity percentage* refers to the active [non-idle] clocked state² within a time window. Load and utilization, used interchangeably here, refer to the same attribute.

- *Scalable activity* refers to the part (or percentage) of the operation whose execution-time scales inversely with frequency.
- *Stall in activity* refers to the part (or percentage) of the operation that involves stall during its completion. While a longer delay actually causes C-state demotion, the instruction stalls referred to here are much shorter and occur with the CPU in CO active state.

Consider a frequency governor on a DVFS system, updating frequency at every periodic time window T . Consider the present instant at time t_0 (“NOW”), as depicted in figure 1. Let the just-completed window T of workload have t_{act} (or t_a) as part of its scalable activity (i.e., t_{act} represents the cumulative portion of activity that keeps the processor busy in execution without the need for any dependent delay or stalls). Let t_{stall} (or t_s) be the

FIGURE 1: **SCALABILITY-RELATED DEFINITIONS**

sub-duration depicting the effective stalls experienced (inter-subsystem dependency stalls). Also let T_{off} (or T_o) be the cumulative duration where the DVFS subsystem is in a deeper C-state, which is a significantly lower power than CO state. The fundamental difference between T_{off} and t_{stall} is that, in the latter case the subsystem is still active CO while experiencing very short dependency stalls; whereas in the former case the delays are large enough and put the subsystem into a momentary deeper C-state.

In the given window, the load execution active-time is the inverse of the frequency f .

Or, in general, in any window:

$$t_{act} \propto 1/f$$

Intel productive performance³ register counts the productive (non-stall) counts proportional to t_{act} . Defining scalability S as a ratio of delta in productive-count (P_{perf}) to active-count (A_{perf}), results in the following general equations:

$$S = \frac{\text{delta } P_{perf}}{\text{delta } A_{perf}} = \frac{t_{act}}{(t_{act} + t_{stall})}$$

Also, the load l (CO percentage) in this window can be defined as:

$$l = \frac{(t_{act} + t_{stall})}{(t_{act} + t_{stall} + t_{off})}$$

Causality

Consider one specific window T between t_{-1} and t_{θ} . Here, the governor has dispensed a frequency f_{θ} . If the governor had dispensed a different frequency f'_{θ} in the *same* window,

the scalable part of the work (i.e., T_a) would have been, say t'_a and thereby t_θ to t'_θ . For practical purposes, the stall cycles can be assumed to be unrelated to the selected frequency, since the stalls are not explicitly dependent on a local subsystem's clock frequency. Hence t_s remain independent as far as their causality with frequency is considered. Similarly, within the same window the number of productive (stall-free) instructions to be executed remains fixed. Specifically, we can generalize and imply that:
Stall time and productive cycle count are invariant of frequency causality.

Note also that frequency and load are causality considerations within a time window, not across. In the rest of the derivation, a reference to change in frequency is not along the time flow, but another possible causality if the frequency were some other.

UTILIZATION [LOAD] TO FREQUENCY CHANGE ESTIMATION

An important aspect in scalability of workload relates to the extent of change in the load caused by a change in frequency.

As shown in figure 2, the sampled delta values of P_{perf} and A_{perf} as ΔP_θ , ΔA_θ within windows ΔTSC results in the equation: can be associated with active, stall, and off time within T:

$$\begin{aligned} t_a &= \frac{1}{f_0} \Delta P_0 \\ t_a + t_s &= \frac{1}{f_0} \Delta A_0 \\ t_a + t_s + t_{off} &= \frac{1}{F_{tsc}} \Delta TSC \end{aligned}$$

2

FIGURE 2: PERF COUNTER DELTA RELATING TO REFERENCE TIME WINDOW

Now, assume the DVFS governor *had* chosen some target f'_θ instead of initial actual frequency f_θ . The requirement is to find a relation that accurately represents the changed final target load l'_θ as a result of this causality. Hence derive estimated load change l_i at time t_i . Let the A_{perf} count in the causality window with frequency f_θ hypothetically transition from the present ΔA_θ to $\Delta A'_\theta$

$$\Delta A'_\theta = f'_\theta(t'_a + t'_s)$$

The stall time is not an explicit function of a local subsystem clock, the time spent in stall remains the same. In other words, *stall time is invariant of frequency causality*.

$$\underline{t'_s} = \underline{t_s}$$

$$\Delta A'_0 = f'_0 \left\{ \frac{\Delta P'_0}{f'_0} + \frac{\Delta A_0 - \Delta P_0}{f_0} \right\}$$

When we solve this with the following two equations for load and scale-factor in general, at time t_n :

$$\Delta A = f_n l_n \Delta TSC / F_{TSC}$$

and the scale factor as

$$s_n = \frac{\Delta P_n}{\Delta A_n}$$

solving yields

$$l'_0 = s'_0 l'_0 + l_0 - s_0 l_0$$

or

$$l'_0 = l_0 \frac{(1 - s'_0)}{(1 - s_0)}$$

Now,

$$s'_0 = \frac{t'_a}{(t'_a + t'_s)}$$

Again, due to stall time invariance,

$$\frac{\left(\frac{1}{s'_0} - 1\right) \Delta P'_0}{f'_0} = \frac{\left(\frac{1}{s_0} - 1\right) \Delta P_0}{f_0}$$

Since this is not an iteration to the next time window, the productive cycles to be executed remain the same. In other words,

The productive cycle count is invariant of frequency causality.

$$\Delta P'_0 = \Delta P_0$$

This results in, the *scaling of scale factor* equation

$$s'_0 = \frac{1}{1 + \frac{f'_0}{f_0} \left(\frac{1}{s_0} - 1 \right)}$$

This intermediate equation defines the changes [scaling] of the scale factor itself over frequency. Solving further, we get the *scaling of load* equation:

$$l'_0 = l_0 \left(s_0 \cdot \frac{f_0}{f'_0} + (1 - s_0) \right)$$

In the causality-based derivation above, there was no workload inherent change in that window. In other words, if the causality from f_θ to f'_θ were actual changes f_1 in adjacent future time window t_1 , and the equation above were used to estimate such a target load l_1 , then any divergence $(L_{actual} - l'_\theta)$ from expected target load is clearly workload inherent change culminated over the time window t_θ to t_1 . Since L_{actual} can be known, any divergence from the estimated load can be calculated. The estimated scaled load is this given by:

$$l_1 = l_0 \left(s_0 \cdot \frac{f_0}{f_1} + (1 - s_0) \right)$$

Similarly, the estimated scaled scale-factor is approximately given by:

$$s_1 = \frac{1}{1 + \frac{f_1}{f_0} \left(\frac{1}{s_0} - 1 \right)}$$

VERIFICATION ON AN ACTUAL PLATFORM

The last two equations above can be called the workload scaling equations. Verifying them empirically on an actual platform, however, isn't as straightforward. At the granularity of a single cycle, the three parameters on the right $[f_\theta, f_1, s_\theta]$ may be determined; as described, the math result cannot be compared directly to L after the end of that cycle; primarily because the incoming load itself could fluctuate as seen at a discrete sample basis. It is possible though, to statistically verify by applying a series of histograms at different frequencies and tracing its ridge. Using this *histogram ridge trace* method, the

FIGURE 3: **A BROWSER-BASED WORKLOAD FROM WEBGLSAMPLES.ORG**

4

generality of the equation is verified.

To begin, a workload is run completely at a fixed frequency.

Workload #1 is a random browser-based graphics load (figure 3; <https://webglsamples.org/multiple-views/multiple-views.html>)

While capturing the data, it is preferred to stop unwanted background services that are irrelevant to the workload under analysis. Also, a precise and light-weight data logging tool is preferred. On Linux systems, the Intel PSST⁴ can accomplish logging at fixed, low, and precise overhead. The data is captured for, say, a duration of 100 seconds and a sampling poll period of 20 ms. Data captured with every sample includes:

FIGURE 4: **UTILIZATION VARIATION IN TIME DOMAIN**

5

FIGURE 5: UTILIZATION HISTOGRAM PLOT AT FIXED FREQUENCY 1 GHz

- CPU load.
- scale factor.
- clamped frequency value.

The entire run is repeated at every frequency of CPU, while keeping other DVFS as fixed sources [e.g., Gfx frequency at a moderate but fixed value].

The CPU scale-factor and utilization logs are captured. In this example, time-domain analysis is not practical for gathering scalability and utilization. Instead, a histogram analysis identifies in which bucket the maximum number of load points lie for a given frequency of operation. The

6

FIGURE 6: COLLECTION OF UTILIZATION HISTOGRAMS FOR EACH FREQUENCY

FIGURE 7: SCALE-FACTOR HISTOGRAMS FOR EACH FREQUENCY

7

8

FIGURE 8: SCALE-FACTOR RIDGE POINTS CONTOUR COMPARED WITH EQUATION

peak of the histogram represents the load value with the highest probability of occurrence for the given frequency. In the example data in figure 5, about 500 samples (or 80 percent of total samples) occurred at a 33 percent load. Such inferences are obviously hard to derive from the time chart in figure 4 even if the frequency is fixed.

The experiment is iterated through all other frequencies, one at a time, to get a range of peak histogram values, as shown in figure 6.

With the same data from the set of experiments above, a similar histogram ridge trace for scale-factor can be plotted as shown in figure 7.

Scaling of scale factor

A plotting tool is used to create a map-view of peak histogram points (in scale factor) in figure 7, which is then

9

FIGURE 9: UTILIZATION RIDGE POINT CONTOUR COMPARED WITH EQUATION

plotted with red dots in figure 8. On the same map-view, the *scaling of scale factor* equation is plotted for different possible target frequencies based on any one initial *seed* value for utilization and scale factor. The correlation between the equation and the full data set is shown in Figure 8.

Scaling of load

The map-view of peak histogram points of utilization (figure 6) is plotted with red dots in figure 9. On the same map view, the *scaling of load* equation is plotted for different possible target frequencies based on one initial seed value for utilization and scale factor.

As summarized by the results shown in these figures, the independent estimates of the equation fit nearly

Related articles

Power-Efficient Software

Eric Saxe

Power-manageable hardware can help save energy, but what can software developers do to address the problem?
<https://queue.acm.org/detail.cfm?id=1698225>

Maximizing Power Efficiency with Asymmetric Multicore Systems

Alexandra Fedorova, et al.

Asymmetric multicore systems promise to use a lot less energy than conventional symmetric processors. How can we develop software that makes the most out of this potential?
<https://queue.acm.org/detail.cfm?id=1658422>

Learning from THE WEB

Adam Bosworth

<http://queue.acm.org/detail.cfm?id=1103833>

precisely with the actual experiment's statistical results. This methodology is applied to multiple other workloads (with different scale factors and load levels) and verified as above.

CONCLUSION

This article elaborated on the causalities of utilization and scale factor in relation to possible frequency-selection choices within a time window. Based on this, generic workload scaling equations are derived, quantifying utilization impact. The equations are verified by applying the histogram ridge trace method at discrete DVFS block level. Though only the CPU core example was detailed, similar equation agreements

were observed on other DVFS blocks (graphics sub blocks), other workloads and other operating systems (Windows®/ Linux®). The equation-estimated curve correlates very accurately with the actual ridge trace curve. This implies that the utilization impact can be “predicted” to be statistically accurate in every cycle and any divergence from it can be attributed to workload-inherent utilization change in that cycle and treated as appropriate to the solution using it.

Acknowledgments

The author is thankful for all the help from colleagues at Intel—specifically, valuable input from principal engineers Harinarayanan Seshadri, and Rajeev Muralidhar and software engineer B. M. Shravan K.

References

1. Howse, B. 2015. Examining Intel's new Speed Shift tech on Skylake: more responsive processors. AnandTech; <http://www.anandtech.com/show/9751/examining-intel-skylake-speed-shift-more-responsive-processors>.
2. Kidd, T. 2014. Power management states: P-states, C-states, and Package C-states. Intel Developer Zone; <https://software.intel.com/en-us/articles/power-management-states-p-states-c-states-and-package-c-states>.
3. Intel. Intel 64 and IA-32 Architectures, Software Developer's Manual, Volume 3B: System Programming Guide, Part 2; <https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-vol-3b-part-2-manual.pdf>.
4. Power Shaping and Stress Tool (PSST) for Intel Platforms. <https://github.com/intel/psst>.

Noor Mubeen is a software architect at Intel client R&D group, focusing on power, thermal and energy management breakthroughs. In the past, he has architected thermal management on Intel mobile devices. He has over 17 years of experience spanning a breadth of domains including networking, storage file systems, and embedded devices. He is an open-source GNU/Linux enthusiast.

Copyright © 2018 held by owner/author. Publication rights licensed to ACM.