Log^2: An Overhead-Constrained Logging Paradigm for Performance Diagnosis

1. Mechanism of Threshold Updating

Modeling. We model the whole threshold adjusting process by three factors: workload $L(t)$, probability distribution of utility score $f(\theta_t, em)$, and utility threshold g. Here t is time, θ_t is some unknown parameter vector to model the probability distribution of utility score, and em is utility score. Figure 1 illustrates an approximate distribution of execution time (we use execution time or latency transparently in this document) of all the MCR (Monitored Code Region). If we choose the formula of utility score as *t*, then the distribution in Figure 1 is $f(\theta_t, em)$.

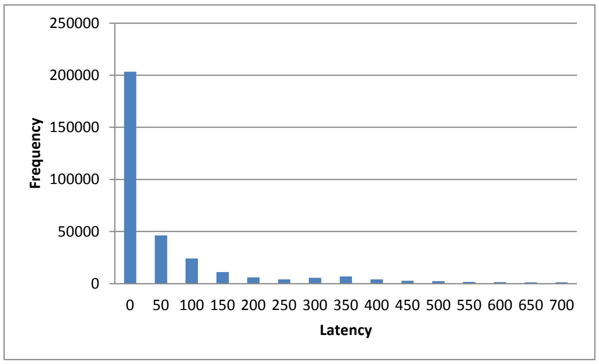

Figure 1. The distribution of the latency of all the LSRs

The overall characteristics of the system execution, such as the behavior of performance problems, latency distribution as well as environment information are all considered, and be abstracted into $L(t)$ and $f(\theta_t, em)$.

Considering the system is slow time varying, i.e., $L(t)$ or $f(\theta_t, em)$ varies slowly in the adjacent update intervals. This assumption is fairly reasonable, since we choose update interval to 30 seconds. And in practice, the environmental dynamics do have negligible changes in such as time scale. Such phenomenon can be represented as follows:

$$\frac{\partial L}{\partial t} \approx 0, \frac{\partial f(\theta_t, x)}{\partial t} \approx 0$$

Given a threshold g, the volume of log instances $b(g,t)$ that inserted into swap buffer, is (without loss of generality, we use continuous form for illustration for simplicity illustration)

$$b(g,t) \\ = \{all\ the\ logs\ with\ utilitys\ score\ larger\ than\ g\} \\ = L(t) \int_g^\infty f(\theta_t, x)dx$$

Updating mechanism. Our goal is to find a g^*, so that the volume is just the budget B (Specifically, if the threshold is set too low, massive logs are inserted into the swap buffer, the consequence is significant memory usage as well as lock contention caused by high-concurrency insertion. In addition, most of these logs are still discarded before flushing. On the other hand, if the threshold is set too high, and thus almost no logs are cached in the buffer, then important logs could be missed, leading to unacceptable logging effectiveness. In summary, the optimal objective is to cache just budget-volume logs), the equation is

$$B = L(t) \int_{g^*}^\infty f(\theta_t, x)dx$$

Additional mathematical deductions are needed. According to the theory of total differential:

$$\Delta b(g,t) \approx \frac{\partial b}{\partial g}\Delta g + \frac{\partial b}{\partial t}\Delta t$$

Here

$$\frac{\partial b}{\partial t} = \frac{\partial L}{\partial t}\int_g^\infty f(\theta_t, x)dx + L(t)\int_g^\infty \frac{\partial f}{\partial t}dx \approx 0$$

Due to the slow time varying property illustrate previously: $\frac{\partial L}{\partial t} \approx 0, \frac{\partial f(\theta_t, x)}{\partial t} \approx 0$

So $\Delta b(g,t) \approx \frac{\partial b}{\partial g}\Delta g$ implies that in a small neighborhood of time t, $b(g,t) \approx b(g)$.

Our objective is now simplified as find the root g^* of function

$$F(g) = b(g) - B$$

We adopt the Secant Method to solve this problem in a numerical manner, specifically

$$g_n = g_{n-1} - F(g_{n-1})\frac{g_{n-1} - g_{n-2}}{F(g_{n-1}) - F(g_{n-2})}$$

Note that g_i is the threshold in the i-th round, and $F(i)$ is the difference or error observed by using g_i in the i-th round.

Intuitively, If error is 0, then g is achieved our goal, we have no change. Otherwise, the offset is proportional to the error multiplied by $\frac{g_{n-1}-g_{n-2}}{F(g_{n-1})-F(g_{n-2})}$, the latter formula is used to approximate the derivative F'. According to the mathematical analysis, the rate of convergence of secant method is approximately 1.618, which is super-linearly (i.e., converges very quickly).

2. Implementation Details

In the threshold adjusting stage, we can change the target B slightly large than actual budget. Such modification brings very small additional overhead, but could mitigate the chance that the volume after local filter is smaller than budget. So the final effectiveness can be better preserved.

To avoid dividing 0 in the formula $\frac{g_{n-1}-g_{n-2}}{F(g_{n-1})-F(g_{n-2})}$, we add the following modification for robustness concern. For the denominator, if $F(g_{n-1}) - F(g_{n-2})$ is very close to 0, we add (if $F(g_{n-1}) - F(g_{n-2}) \geq 0$) 1or -1 (if $F(g_{n-1}) - F(g_{n-2}) < 0$) to it. when $g_{n-1} - g_{n-2}$ equals to zero, threshold will trap to a certain number and never changes regardless to any dynamics. To avoid this, a very small value 0.01 is added under such condition. Note that the choice of 0.01 or 1 is relatively flexible, varying it to other small values will not bring significant difference.