Optimal quorum for a reliable Desktop grid

Ilya Chernov and Natalia Nikitina

Inst. Applied Math Research, Pushkinskaya 11, Petrozavodsk, Russia

Abstract. We consider a reliable desktop grid solving multiple tasks with answers from a given set. Wrong answer can be obtained with some small probability p: reliability means that p is small with respect to $1-p$. Penalty is added to the computation cost if a wrong answer is believed to. We consider the optimization problem of choosing the quorum in case of finite set of possible answers, countable set, and set with continuously distributed subset.

Keywords: Desktop grid, optimal quorum, replication, reliable computation

1 Introduction

In the presented work we consider the problem of task scheduling in a Desktop Grid. The term stands for a distributed computing system with computational nodes voluntarily donated by an organization (or a group of organizations following the same research goals) in their idle time. The nodes of such system are desktop PCs, compute servers, cluster nodes, and other computational resources available in a local network and/or the Internet. The BOINC middleware [1] can be considered a de-facto standard for organizing Desktop Grids. Since late 1990s, BOINC-based Desktop Grids have proven to serve as an effective, highly scalable tool for solving computationally intensive problems. But the definition given above implies that computational nodes of a Desktop Grid can be highly heterogeneous by technical and software characteristics and become available/unavailable in unpredicted moments. Moreover, the answer returned by the computing node can be wrong, due to malfunction, errors in transaction, malicious actions etc. In order to harness the volatile Desktop Grid resources rationally, special scheduling methods and algorithms should be developed. Much effort has been paid recently to optimize the calculation process in a Desktop Grid and improve its reliability. For example, a series of works [3, 4, 6] consider heuristics for mapping independent tasks onto identical nodes that may leave Desktop Grid in random moments; the work [5] presents closed-form conditions of task replication to be reasonable; the authors of [2] and [7] consider heterogeneity of computational nodes, etc. The contribution of the presented work is investigation of the optimal task replication level and the optimal quorum for a Desktop Grid with high reliability of the answers returned from computational nodes. Solving a computational problem with high reliability may significantly increase the cost of computations, such as the overall time, due to high precision of calculations. We try to evaluate the expected cost and select the optimal strategy of task scheduling.

2 The mathematical model

Let us consider a desktop grid computing system that consists of multiple computers and solves numerous similar tasks. Each task has an answer from some set of possible answers. Without loss of generality we can assume that this set is a subset of real numbers or integer numbers. The correct answer is obtained with some probability q. Errors are possible due to malfunction of hardware, malicious actions, or wrong answers produced by a correct non-deterministic algorithm. Other, wrong answers have the total probability $1 - q$; some of them can have non-zero probability, others are distributed continuously.

Such errors can be made less likely by replication: sending copies of each task to different computing nodes until a chosen number ν of identical answers is received. This number is called the quorum. Obviously, for a given quorum the number of copies is at least ν but can be more, up to $N(\nu - 1) + 1$ for finite number N of possible answers. However, the average number of copies is finite even in case of infinite set of answers.

If a wrong answers is taken, this error will later be revealed and can cost much. The losses can be connected with loosing a rare desired phenomenon, unnecessary expensive laboratory checks, reputational losses, etc. These penalties can be huge compared to the average cost of a single task.

Here we assume that $q \approx 1$ in the following sense: $1 - q$ is negligibly small compared to q. Therefore we can neglect possible cases of choosing between different wrong answers or any wrong answers seen if the correct one has been believed.

First we consider the finite set of possible answers. Then we will add some simple notes about more general case with infinite set of answers, with positive probabilities and/or continuously distributed.

3 Finite set of answers

Assume that $M + 1$ possible answers have probabilities p_m, $p_0 = q \approx 1$ (the correct answer), penalties of different answers are F_m, $F_0 = 0$. If quorum ν is chosen and the correct answer is accepted, the average number of copies is ν up to the precision discussed above, the probability to get the correct answer is q. However, for each wrong answer number m possible durations can be any between ν (the same wrong answer in a row) and $2\nu - 1$ (after $\nu - 1$ correct and $\nu - 1$ wrong answers the final wrong answer was received). The corresponding probabilities of durations $\nu + i$, $i = 0, \dots, \nu - 1$ are, respectively, equal to

$$\binom{\nu - 1 + i}{i} p_m^\nu q^i \approx \binom{\nu - 1 + i}{\nu - 1} p_m^\nu.$$

The first factor chooses i positions of correct answers among $\nu - 1$ first tries, while the final answer is accepted.

The mean cost is

$$E(\nu) = \nu + \sum_{m=1}^{M} F_m \sum_{i=0}^{\nu-1} \binom{\nu-1+i}{\nu-1} p_m^{\nu}.$$

To simplify this expression, we need the following formula:

$$\sum_{i=0}^{\nu-1} \binom{\nu-1+i}{\nu-1} = \binom{2\nu-1}{\nu}.$$

It follows from

$$\sum_{j=1}^{\nu} \binom{(2\nu-1)-j}{\nu-1} = \binom{2\nu-1}{\nu-1}$$

if $j = \nu - i$. To prove this equality, we need two well-known formulae (the Pascal triangle):

$$\binom{n}{m} = \binom{n-1}{m-1} + \binom{n-1}{m}, \quad \binom{n}{m} = \binom{n}{n-m}.$$

Now consider

$$\binom{2\nu-1}{\nu-1} = \binom{2\nu-2}{\nu-2} + \binom{2\nu-2}{\nu-1} = \binom{2\nu-2}{\nu-1} + \binom{2\nu-2}{\nu} =$$
$$= \binom{2\nu-2}{\nu-1} + \binom{2\nu-3}{\nu-1} + \binom{2\nu-3}{\nu}.$$

Continue to obtain

$$\binom{2\nu-1}{\nu-1} = \sum_{j=1}^{\nu} \binom{(2\nu-1)-j}{\nu-1}.$$

On the final step we applied the obvious equality $\binom{n}{n} = \binom{n-1}{n-1}$. This completes the proof.

Now we can rewrite the mean cost as

$$E(\nu) = \nu + \binom{2\nu-1}{\nu} \sum_{m=1}^{M} F_m p_m^{\nu}. \tag{1}$$

We want to minimize this cost keeping in mind the fact that penalties F_m are not known precisely, so that conditions should allow variations of F_m. Let us consider differences $\Delta E_\nu = E(\nu+1) - E(\nu)$ and take the first ν with $\Delta E_\nu \geq 0$. It is easy to check that

$$\Delta E_\nu \approx 1 - \binom{2\nu-1}{\nu} \sum_{m=1}^{M} F_m p_m^{\nu}.$$

So we need the smallest ν such that

$$\binom{2\nu-1}{\nu}\sum_{m=1}^{M} F_m p_m^{\nu} \leq 1. \tag{2}$$

It is clear that inequality (2) holds, then each term is less than 1; although it is possible to choose parameters in such a way that all terms are less than one while the sum is more. However, if we take ν making each term less than 1, difference between optimal value is at most 1.

Let us simplify condition (2) in case $M = 1$ using the asymptotical formulae

$$\binom{2\nu}{\nu} = \frac{2^{2\nu}}{\sqrt{\pi\nu}}, \quad \binom{2\nu-1}{\nu} = \frac{2^{2\nu-1}}{\sqrt{\pi\nu}}.$$

Note that they are rather precise even for low ν: 10% for $\nu = 1$ and less. Then the condition (2) can be replaced by the approximate one:

$$\frac{2^{2\nu-1}}{\sqrt{\pi\nu}} F p^{\nu} \leq 1.$$

If $F = 2^A$, $p = 2^{-B}$ ($A > 1$, $B > 2$), then

$$(2 - B)\nu + A - 1 \leq 0.5 \log_2 \pi\nu.$$

The right-hand side is less than 3 for $\nu < 20$ and we can safely neglect it; then

$$\nu \geq \frac{A-1}{B-2} = -\frac{\log_2(F) - 1}{\log_2(p) + 2} = -\frac{\ln(F/2)}{\ln(4p)},$$

the smallest ν to be taken. For example, for $A = 20$, $B = 10$ we have $\nu = 3$. The same quorum we get solving the general inequality (2).

4 Infinite set of wrong answers

Assume that beside the correct answer of some probability $q \approx 1$ and M wrong ones with probabilities $p_m > 0$ the set of possible answers contains a subset S of continuously distributed answers of total probability r. Obviously these answers are thrown off by any replication because they have zero probabilities and therefore receiving an answer again is an impossible event. However, wrong answers with positive probabilities can come a few times and be believed. In this case the cost is increased on some penalty F_m.

The results here are completely the same as in the previous section, only the average cost of $\nu = 1$ case is different because estimated value of penalty for an answer from S must be added to the average cost.

Now let us consider the case of countable ($M = \infty$) set of answers with positive probabilities. This assumption may look nice, for example in case when answers are integer numbers. However, such case does not allow equal (or at least comparable) probabilities p_m: the series

$$\sum_{m=1}^{\infty} p_m$$

must converge to some p which is small compared to $1 - p$. Provided that p_m are given, the average cost

$$E(\nu) = \nu + \sum_{m=1}^{\infty} \binom{2\nu - 1}{\nu} F_m p_m^{\nu}$$

can still diverge for some or all ν, depending on behaviour of penalties F_m. It is clear that polynomial (at most) growth of F_m with respect to m is sufficient for convergence. To prove that, note that a binomial coefficient grows with polynomial rate:

$$\binom{n}{m} \le 2^n, \quad 0 \le m \le n,$$

and the sequence p_m decreases more quickly that m^{-1} because it forms a convergent series; therefore sufficiently large $\nu = \nu^*$ makes the terms of the series decrease quickly enough, so the the series converges. Higher ν only reduce the sum at least $(\max p_m)^{\nu-\nu^*}$ times, so that the second term in the expression for E decreases. So, for high enough values of ν the expected cost E grows, almost lineary. Then it has the unique minimal value being the solution to the optimization problem. We have proven that the optimization problem we are studying always has a solution. It is solved in the similar way as the one for the finite set of answers, only the cost function may be equal to infinity for some μ. Note that the solution ν can equal 1 in case of low (or no) penalties. This means that no replication is necessary.

It is clear that both countable set of possible answers and continuously distributed set of impossible answers can also be combined.

Acknowledgements

The work was financially supported by grants 13-07-00008 and 15-29-07974 of Russian Foundation for Basic Research.

References

1. D. Anderson. BOINC: A system for public-resource computing and storage. In *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*, pages 4–10, Washington, DC, USA, 2004.
2. C. Anglano, J. Brevik, M. Canonico, D. Nurmi, and R. Wolski. Fault-aware scheduling for bag-of-tasks applications on desktop grids. In *Proceedings of the 7th IEEE/ACM International Conference on Grid Computing*, pages 56–63, 2006.

3. L.-C. Canon, A. Essafi, G. Mounie, and D. Trystram. A bi-objective scheduling algorithm for desktop grids with uncertain resource availabilities. In *Proceedings of the 17th International Conference on Parallel Processing (Euro-Par11)*, volume II, pages 238–249, Berlin, Heidelberg, 2011. Springer-Verlag.
4. H.-C. Hwang, K. Lee, and S. Y. Chang. The effect of machine availability on the worst-case performance of LPT. *Discrete Applied Mathematics*, 148(1):49–61, 2005.
5. O. Ibarra and C. Kim. Heuristic algorithms for scheduling independent tasks on nonidentical processors. *Journal of ACM*, 24(2):280–289, 1977.
6. C.-Y. Lee. Parallel machines scheduling with nonsimultaneous machine available time. *Discrete Applied Mathematics*, 30(1):53–61, 1991.
7. A. S. Rumiantsev. Optimizing the execution time of a desktop grid project. *Program Systems: Theory and Applications online journal*, 5(1(19)):175–182, 2014. (In Russian).