A lower limit for the probability of success of computing tasks in a grid

Christer Carlsson

IAMSR, Åbo Akademi University,
Joukahaisenkatu 3-5 A, FIN-20520 Turku
e-mail: christer.carlsson@abo.fi

Robert Fullér

IAMSR, Åbo Akademi University,
Joukahaisenkatu 3-5 A, FIN-20520 Turku
e-mail: rfuller@abo.fi

József Mezei

Turku Centre for Computer Science,
Joukahaisenkatu 3-5 B, FIN-20520 Turku
e-mail: jmezei@abo.fi

Abstract: We will consider an approximation to the probability that a particular computing task in a grid is successful. This happens if there will always be at least a single idle node available in the system in the case of a node failure. In this paper we will show a lower limit for the probability of success of a computing task in a grid.

1 Introduction

Grid computing is a form of distributed computing where clusters of networked, loosely-coupled computers, are acting in concert to perform very large and/or a large number of tasks [6]. However, running applications on the Grid environment poses signicant challenges due to the diverse failures encountered during execution. If one computing node fails during the job execution, the whole job fails and has to be restarted. To handle resource failures and avoid restarting the job from the initial state, the fault-tolerance mechanisms checkpointing and migration have been developed. The Resource Management System makes a snapshot of the jobs execution state, transfers the snapshot to another computing node, and resumes the job execution after all nodes failed have been replaced with nodes that are working [8]. Users negotiate for resource usage through a Grid resource broker which queries resource providers on their behalf to find suitable resources. They require a job execution with a desired level of priority and quality [5]. Failure intensity usually increases with age for mechanical equipment. Power law and loglinear Poisson processes are often used to model failure intensity. The distinguishing feature of Poisson processes is that the previous history of failure times $t_1, \ldots, t_n$ does not affect failure intensity [1]. We consider an approximation to the probability that a particular computing task in a grid is successful. This happens if there will always be at least a single idle node available in the system in the case of a node failure. Let *success* denote the event that the task is successful, and *failure* the opposite event. We formulate the probability of the success as the sum of the probabilities $P(\text{"none of the nodes allocated to the task fail"}) + \sum_{m=1}^{m_{\max}} P(\text{"}m \text{ of the nodes allocated to the task fail \& at least } m \text{ idle nodes are available as reserves"})$. Here $m_{\max}$ is an upper limit for the number of failures considered. The value can be chosen by judging the size of the contribution of each event, determined by the corresponding probability. Thus, the sum can be simplified by considering only those events that do not have vanishingly small probabilities. A conservative bound for the success probability can be derived by assuming that the m failures take place simultaneously, which leads to

$$
\begin{aligned}
P(\text{success}) &= 1 - P(\text{failure}) \\
&= 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur \& less than } m \text{ free nodes available}) \\
&= 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) P(\text{less than } m \text{ free nodes available}) \\
&\geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) P(\text{less-than-m-anytime})
\end{aligned}
$$

Here *less-than-m-anytime* stands for the event *less than m free nodes available at any time point*.

2 Copulas

Copulas characterize the relationship between a multidimensional probability function and its lower dimensional margins. A two-dimensional copula is a function $C\colon [0,1]^2 \rightarrow [0,1]$ which satisfies

- $C(0, t) = C(t, 0) = 0$
- $C(1, t) = C(t, 1) = t$ for all $t \in [0, 1]$
- $C(u_2, v_2) - C(u_1, v_2) - C(u_2, v_1) + C(u_1, v_1) \geq 0$ for all $u_1, u_2, v_1, v_2 \in [0, 1]$ such that $u_1 \leq u_2$ and $v_1 \leq v_2$.

Equivalently, a copula is the restriction to $[0, 1]^2$ of a bivariate distribution function whose margins are uniform on $[0, 1]$. The importance of copulas in statistics is described in the following theorem,

Theorem 2.1 (Sklar, [9]). *Let X and Y be random variables with joint distribution function H and marginal distribution functions F and G, respectively. Then there exists a copula C such that*

$$H(x, y) = C(F(x), G(y)),$$

Moreover, if marginal distributions, say, $F(x)$ and $G(y)$, are continuous, the copula function C is unique. Conversely, for any univariate distribution functions F and G and any copula C, the function H is a two-dimensional distribution function with marginals F and G.

Thus copulas link joint distribution functions to their one-dimensional margins and they provide a natural setting for the study of properties of distributions with fixed margins. For further details, see Nelsen [7]. For any copula C we have,

$$W(u, v) = \max\{0, u + v - 1\} \leq C(u, v) \leq M(u, v) = \min\{u, v\}.$$

In the statistical literature, the largest and smallest copulas, M and W are generally referred to as the Fréchet-Hoeffding bounds.

3 A lower limit for probability of success

To simplify the inference about the length of a task affecting a number of nodes we assume that the length follows a Gaussian distribution. Let us introduce the notations $F(t) = P$(the length of a task is less than t) and let t^* be chosen such that $1 - F(t^*) \geq 0.995$. Let the data $t_1, ..., t_b$ represent the lengths (say, in minutes) of b tasks. This leads to the sample mean $\bar{t} = \sum_{i=1}^{b} t_i$ and variance $s^2 = b^{-1} \sum_{i=1}^{b} (t_i - \bar{t})^2$. Assuming the standard reference prior (see [3]) for the parameters, we obtain the predictive distribution for the length of a future task, say

T, which has the T-distribution with parameters $\bar{t}, ((b-1)/(b+1))s^2, b-1$, i.e. the probability density of the distribution equals

$$p(t|\bar{t}, ((b-1)/(b+1))s^2, b-1) = \frac{\Gamma(b/2)}{\Gamma(\frac{b-1}{2})\Gamma(1/2)} \left(\frac{1}{(b+1)s^2}\right)^{\frac{1}{2}} \times \left[1 + \frac{1}{(b+1)s^2}(t-\bar{t})^2\right]^{-\frac{b}{2}}. \quad (1)$$

Furthermore, let

$$G(m) = P(\text{less-than-m-anytime})$$

and let us denote the copula of F and G by H, where

$$H(t, m) = P(\text{the length of a task is less than t, less-than-m-anytime}),$$

then we have

$$H(j, m) = C(F(j), G(m) \leq M(F(j), G(m) = \min(F(j), G(m)) \quad (2)$$

Using theses notations we find,

$$P(\text{success}) = 1 - P(\text{failure}) = 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur \& less-than-m-anytime})$$

$$= 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) P(\text{less than } m \text{ free nodes available})$$

$$\geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) P(\text{less-than-m-anytime})$$

$$\geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur})$$

$$\times \left\{ \sum_{j=1}^{t^*} P(j-1 \leq \text{the length of a task} < j, \text{less-than-m-anytime}) \right.$$

$$\left. + P(\text{the length of a task} \geq t^*, \text{less-than-m-anytime}) \right\}$$

$$\geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur})$$

$$\times \left\{ \sum_{j=1}^{t^*} P(\text{the length of a task} < j, \text{less-than-m-anytime}) \right.$$

$$\left. + P(\text{the length of a task} \geq t^*, \text{less-than-m-anytime}) \right\}$$

Finally, using relation (2) we get,

$$
\begin{aligned}
P(\text{success}) \geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) \\
\times \Bigg\{ \sum_{j=1}^{t^*} \min \Bigg\{ \frac{\Gamma(b/2)}{\Gamma(\frac{b-1}{2})\Gamma(1/2)} \left(\frac{1}{(b+1)s^2} \right)^{\frac{1}{2}} \left[1 + \frac{1}{(b+1)s^2} (j - \bar{j})^2 \right]^{-\frac{b}{2}}, \\
P(\text{less than } m \text{ free nodes at any time point}) \Bigg\} \\
+ \underbrace{P(\text{the length of a task is } \geq t^*, \text{less-than-m-anytime})}_{\leq 0.005} \Bigg\}
\end{aligned}
$$

Summarizing our findings we get,

$$
\begin{aligned}
P(\text{success}) \geq 1 - \sum_{m=1}^{m_{\max}} P(m \text{ failures occur}) \\
\times \Bigg\{ \sum_{j=1}^{t^*} \min \Bigg\{ \frac{\Gamma(b/2)}{\Gamma(\frac{b-1}{2})\Gamma(1/2)} \left(\frac{1}{(b+1)s^2} \right)^{\frac{1}{2}} \left[1 + \frac{1}{(b+1)s^2} (j - \bar{j})^2 \right]^{-\frac{b}{2}}, \\
P(\text{less-than-m-anytime}) \Bigg\}
\end{aligned}
$$

4 Summary

Using the Fréchet-Hoeffding bounds for copulas we have shown a lower limit for the probability of success of a computing task in a grid.

5 Acknowledgment

This work has been partially supported by the EU within the 6th Framework Programme under contract IST-031772 *Advanced Risk Assessment and Management for Trustable Grids*, (AssessGrid).

References

[1] R.D. Baker: Calculating the expected failure rate of complex equipment subject to hazardous repair, International Journal of Production Economics 67 (2000) 53-61.

[2] C. Baudrit, I. Couso, D. Dubois: Joint propagation of probability and possibility in risk analysis: Towards a formal framework, International Journal of Approximate Reasoning, 45 (2007) 82-105.

[3] J.M. Bernardo, A.F.M. Smith, *Bayesian Theory* (Wiley, 1994).

[4] M. Boullón et al., Modeling Execution Time of Selected Computation and Communication Kernels on Grids, in: Advances in Grid Computing - EGC 2005, Lecture Notes in Computer Science, vol. 3470, 2005 731-740.

[5] K. Djemame, I. Gourlay, J. Padgett, G. Birkenheuer, M.Hovestadt, O. Kao, K. Voss: Introducing Risk Management into the Grid, in: Proceedings of the Second IEEE International Conference on e-Science and Grid Computing, 2006 28-35.

[6] I. Foster: What is the Grid? A Three Point Checklist, GRIDToday, 2002.

[7] R.B. Nelsen, *An Introduction to Copulas*, (Springer, New York, 1999).

[8] Kerstin Voss: Recursive Evaluation of Fault Tolerance Mechanisms for SLA Management, in: Proceedings of the Fourth International Conference on Networking and Services, 2008 223-229.

[9] A. Sklar, Fonctions de Répartition àn Dimensions et Leurs Marges. Paris, France: Publ. Inst. Statist. Univ. Paris, 1959, vol. 8, pp. 229-231.