A framework for dependable QoS adaptation in probabilistic environments [*]

A. Casimiro
Univ. of Lisboa
casim@di.fc.ul.pt

P. Lollini
Univ. of Firenze
lollini@dsi.unifi.it

M. Dixit
Univ. of Lisboa
mdixit@di.fc.ul.pt

A. Bondavalli
Univ. of Firenze
bondavalli@unifi.it

P.Veríssimo
Univ. of Lisboa
pjv@di.fc.ul.pt

ABSTRACT

Distributed protocols executing in uncertain environments, like the Internet, had better adapt dynamically to environment changes in order to preserve QoS. In a previous work, it was shown that QoS adaptation should be dependable, if correctness of protocol properties is to be maintained. In this paper we provide concrete strategies and methodologies to improve the implementation of dependable QoS adaptation. During its lifetime, a system alternates periods where its temporal behavior is well characterized, with transition periods where a variation of the environment conditions occurs. Our method is based on the following: if the environment is generically characterized in analytical terms, and we can detect the alternation of these stable and transient phases, we can drastically improve the effectiveness of dependable QoS adaptation. To prove our point, we conduct an evaluation based on "synthetic" data flows generated from one or more probabilistic distributions, and we show that the proposed strategies can indeed be effective and still dependable in the considered cases.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems

General Terms

Reliability

Keywords

Dependability, QoS adaptation, probabilistic QoS

1. INTRODUCTION

Computer systems and applications are becoming increasingly distributed and we witness to the pervasiveness and ubiquity of computing devices. This openness and complexity means that the environment tends to be unpredictable, essentially asynchronous, making it impractical, or even incorrect, to assume any time-related bounds. On the other hand, there are increased concerns about the dependability of these systems and applications, in the sense of their ability to meet some specified quality of service (QoS) levels. One possible way to cope with the uncertain timeliness of the environment while meeting dependability constraints consists in building adaptive applications and ensure that they adapt in a dependable way, that is, they remain correct as a result of adaptation.

In this paper we build on previous work that introduced the necessary architectural and functional principles over which a dependable adaptation approach should be based [2]. In essence, the idea behind dependable adaptation is to ensure that a *coverage stability* property is satisfied, that is, that the assumed bounds for fundamental variables are secured with a known and constant probability. Clearly, this is only possible if some assumptions about the environment behavior are made. Then, using appropriate mechanisms and approaches it may be possible to probabilistically characterize the current operational state and derive the bounds that must be used to secure the coverage stability property.

In this paper we introduce a framework that is based on the use of probabilistic methods for the recognition of the "state" of the environment. The framework does not require a specific method to be used, but it allows several methods to be used even simultaneously.

One important and distinguishing aspect of our work is that we are concerned with dependability objectives. We performed a number of simulation experiments using synthetic data flows generated from well-known probabilistic distributions. Based on these results we are able to conclude that the proposed framework may indeed allow to achieve dependable adaptation and improved time bounds, provided that adequate environment recognition methods are used for a given environment behavior.

The paper is organized as follows. Next section provides a motivation for this work and we discuss related work. Then, Section 3 describes the proposed framework for dependable adaptation in probabilistic environments. Implementation issues are then addressed in Section 4, while some simulation results are discussed in Section 5. Conclusions and future perspectives are finally presented in Section 6.

[*]This work was partially supported by the EC, through project IST-FP6-STREP-26979 (HIDENETS), and by the FCT, through the Large-Scale Informatic Systems Laboratory (LaSIGE).

SAC'08 March 16-20, 2008, Fortaleza, Ceará, Brazil
Copyright 2008 ACM 978-1-59593-753-7/08/0003 ...$5.00.

2. MOTIVATION AND RELATED WORK

Providing QoS guarantees for the communication in spite of the uncertain or probabilistic nature of networks is a problem with a wide scope, which can be addressed from many different perspectives. We are fundamentally concerned with timeliness issues and with securing or improving the dependability of adaptive applications.

The fundamental architectural and functional principles for dependable QoS adaptation have been previously introduced in [2]. In this earlier work we also followed a dependable perspective, and analyzed why systems would fail as a result of timing assumptions being violated, as it may happen in asynchronous environments. A relevant effect is *decreased coverage* of some time bound [13], when the number of timing failures goes beyond an assumed limit. This effect can be handled in adaptive systems, by adapting the assumed bounds during the execution to ensure that coverage of that bound is secured. In other words, the objective is to satisfy the so called *Coverage Stability* property. However, deciding when and how to adapt depends on what is assumed about the environment. In our previous work ([2]) a conservative approach was followed, just assuming a probabilistic environment but not a specific probabilistic distribution for delays. Therefore, this led to a conservative solution with respect to the bounds required to guarantee some coverage.

Interestingly, in the last few years several works have addressed the problem of probabilistically characterizing the delays in IP-based networks based on real measured data, allowing to conclude that observed empirical delay distributions may be characterized by well known distributions, such as the Weibull distribution [10, 6], the shifted gamma distribution [11, 4], the exponential distribution [8] or the truncated normal distribution [5]. Based on this, we realized that it would be interesting and appropriate to consider less conservative approaches, by assuming that specific distributions may be identified and thus allowing to achieve better time bounds for the same required coverage.

However, some of these works also recognize that probabilistic distributions may change over time (e.g. [11]), depending on the load or other sporadic occurrences, like failures. Therefore, in order to secure the required dependability attributes, it becomes necessary to detect changes in the distribution and hence use mechanisms for being able to do that. Fortunately, there is also considerable work and well-known approaches addressing this problem (see [14] for a nice overview). Among others, we can find approaches based on time-exponentially weighted moving histograms [9] and on the Kolmogorov-Smirnov test [5].

A fundamental distinguishing factor of our work is that we are concerned with system dependability. Therefore, while other works addressing adaptive systems are concerned with performance improvements, our objective is to achieve dependable adaptive designs by ensuring coverage stability.

3. THE ADAPTATION FRAMEWORK

3.1 Assumptions

As mentioned in the previous section, in this work we advance on previous results by making more optimistic but not less realistic assumptions, in order to achieve improved and still dependable $\langle bound, coverage \rangle$ pairs. Instead of making the weak but restrictive assumption, that the environment is probabilistic but the distribution is unknown (as done in [2]), we now make the following assumptions:

Interleaved known probabilistic behavior - We assume that the environment alternates stable periods, during which it follows some known probabilistic behavior, with unstable periods, during which the probabilistic behavior is unknown or cannot be characterized. As discussed in Section 2, this assumption is supported by the results of many recent works (e.g. [6, 10, 11]).

Recognition abilities - The system has sufficient resources and can be equipped with the adequate mechanisms to detect changes from stable to unstable periods (and vice-versa) and recognize the best fitting probabilistic distribution during stable periods.

Sufficient stability - We assume that probabilistic changes are slower than the detection and recognition speed. In other words, we assume that typical execution environments are sufficiently stable for the time it takes to recognize the actual probabilistic distribution. This is a mandatory assumption for any application that needs to recognize the state of a dynamic environment.

Sufficient activity - We assume that there is sufficient system activity, and therefore enough observation points that we can use to feed some trend detection and recognition mechanisms. Hence it is possible to use measured delays, assuming that all of them report to a sufficiently recent observation period as required for accurately characterizing the state of the environment. Obviously, system activity depends on the application. We believe that this is an acceptable assumption for many systems, such as car-to-car and multimedia applications.

3.2 Environment recognition and adaptation

The proposed framework for dependable (QoS) adaptation can be seen as composed by two activities:

Environment recognition - The environment conditions can be inferred analyzing a real time data flow representing, for example, the end-to-end message delays in a network. Using statistics the data may be described by a probability density function (pdf).We note that the data models can be so complex that they can not be described in terms of simple well-know probabilistic distributions. In the case in which the model that describes the data is known, the problem is reduced to estimating unknown parameters of a known model from the available data. In our experimental implementation of the framework we used the method of maximum likelihood estimation (MLE), which is considered to be one of the most robust techniques for parameter estimation [12].

QoS adaptation - Once the most fitting distribution (together with its parameters) has been identified, its statistics properties can be exploited to find a pair $\langle bound, coverage \rangle$ that will satisfy the objective of keeping a constant coverage of the assumed bound throughout the execution.

3.3 The adaptive approach

As detailed in Section 3.1 we are assuming an interleaved probabilistic behavior of the environment. Therefore, we can consider that the system alternates periods during which the conditions of the environment remain fixed (**stable phases**), with periods during which the environment conditions change (**transient phases**). In the first case, the statistical process that generates the data flow is under control and then we can compute the corresponding distribution using an appropriate number of samples; On the contrary, if the environment conditions are changing, then the associated statistical process is actually varying, so no fixed distribution can describe its real behavior.

During the transient phases we adopt a **conservative approach** and we set the pair $\langle bound, coverage \rangle$ using a probabilistic formulation based on the one-sided inequality [2], which provides a pessimistic bound, but which holds for all probabilistic distributions. As soon as the presence of a stable phase is detected, a proper pdf is identified and then an improved (lower) bound can be computed according to the new distribution (**optimistic approach**), still ensuring the coverage stability property. The bound adaptation is then triggered by the detection of a new stable/transient phase.

3.4 Phase detection mechanisms

A phase detection mechanism must detect the beginning of a new transient phase, as soon as the environment conditions start changing ("changing environment" detection time), as well as the beginning of a new stable phase, as soon as the environment conditions stabilize ("stable environment" detection time).

The "changing environment" detection time is of particular importance in this context, since during such critical periods, the environment is changing but the bound is tailored for a particular set of condition environments that do not hold anymore. The ideal situation arises when the detection of a changing environment is instantaneous.

Although marginal to the main focus of the paper, here we propose some simple phase detection mechanisms that can be implemented in the framework. Such mechanisms can be used individually, or combined together.

- **Mechanism 1**. A stable period with distribution $\hat{D}$ is detected when the estimated mean $E(\hat{D})$ and the estimated variance $V(\hat{D})$ satisfy some relations that are always true for the postulated distribution $\hat{D}$ (e.g., $E(\hat{D})^2 = V(\hat{D})$ for the exponential distribution).

- **Mechanism 2**. At the reception of the k-th sampling point, we compute the expected mean $E(\hat{D})_k$ and observed variance $V(\hat{D})_k$ using the last h collected sampling points, with $k \geq h$. Let $d \in \mathbb{N}^+$ the step-size, a stable phase will be detected if the following f measurements $E(\hat{D})_i$ and $V(\hat{D})_i$, with $i = k + d, k + 2 * d, \ldots, k + f * d$ (each one performed using the available last h collected sampling points), will remain within a given tolerance level e_{dev} and v_{dev} from $E(\hat{D})_k$ and $V(\hat{D})_k$, respectively.

- **Mechanism 3**. A stable period with distribution $\hat{D}$ is detected when some goodness-of-fit (GoF) tests establish the goodness of fit between the postulated distribution $\hat{D}$ and the evidence contained in the experimental observations. For example, the Kolmogorov-Smirnov test [3] is the preferred GoF test in case of a continuous population distribution.

In the following section we describe how these mechanisms have been implemented in the context of the proposed framework. This also serves to clarify the conditions under which the experimental results presented in Section 5 are obtained.

4. IMPLEMENTATION ISSUES

The scheme depicted in Figure 1 shows that the framework can be modelled as a service that: *(i)* accepts the history size (i.e., the number of collected samples of the random variable under observation) and the required coverage as dependability related parameters; *(ii)* reads samples/measured delays as input, using them to fill up the history buffer that is used by the phase detection mechanisms and for the estimation of distribution parameters; and *(iii)* provides, as output, a bound that should be used in order to achieve the specified coverage.

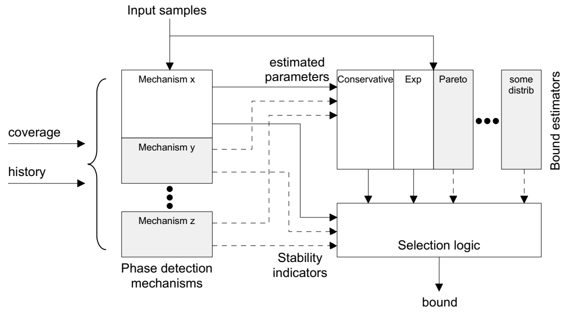

Figure 1: Schematic view of the framework for dependable adaptation.

In this experimental implementation we only considered two probabilistic distributions: the **exponential** and the **Pareto** distributions. This was an early decision that we took based on the fact that these are commonly used distributions and would perfectly serve our purpose of testing the framework. In fact, in a deployed system, the number of different mechanisms that could be used and configured for other, different distributions, would only be limited by the available resources, assumed to be enough by construction, as per the "Recognition abilities" assumption (see Section 3.1).

4.1 Setting up parameters

As previously mentioned, it is possible to use several phase detection mechanisms simultaneously. We have implemented the three mechanisms described earlier to detect the exponential distribution, and "Mechanism 1" for the detection of the Pareto distribution.

"Mechanism 1" detects a stable exponential period if $|\sqrt{V(\hat{D})} - E(\hat{D})| < \tau$, being $E(\hat{D})$ and $V(\hat{D})$ an estimation of the expected mean and variance, respectively, and τ a tolerance bound to encompass for the estimation errors (see [1] for a detailed explanation). The configuration of the other two mechanisms for the exponential distribution was done exactly as explained in Section 3.4, which required some values to be postulated. In particular, for the case of "Mechanism 2" we set the step-size to 1 (thus obtaining the promptest adaptation possible, at the cost of additional processing resources) and, for an history size of $h = 30$, we used $f = 10$. In practical settings these two values must be configured to match the expected environment dynamics. As per the "Sufficient stability" and "Sufficient activity" assumptions, the observed environment will present time windows during which a sufficient number of samples can be observed (at least $h + f$) while the probabilistic behavior is stable. To tune the method sensibility we set $e_{dev} = 0.5 \times E(\hat{D})_k$ and $v_{dev} = 0.5 \times V(\hat{D})_k$. For the case of "Mechanism 3", we just had to postulate the critical values of Kolmogorov-Smirnov statistics for the exponential distribution with unknown mean, which we did by using standard values provided in statistics tables.

With respect to the instantiation of "Mechanism 1" for the Pareto distribution, we had to postulate a value for τ, which in practice depends on the history size and on the shape parameter of the distribution, as shown in Table 1. This table was constructed based on several tests to empirically determine which values would ensure the dependability of the mechanism. The location parameter of the distribution is always set to the lower value observed in the history. A stable Pareto distribution is detected when $|V(\hat{D}) - E(\hat{D})^2/\alpha(\alpha - 2)| < \tau$.

Table 1: Values for τ, Pareto distribution.

α	$\tau(h = 30)$		α	$\tau(h = 30)$
(0,1.2]	1		(2.4,3.2]	0.1
(1.2,2.4]	0.4		(3.2,inf]	0.05

4.2 Defining bound estimators

Depending on the output of the phase detection mechanisms, one of the bounds calculated by the implemented bound estimators will be selected as the output of the framework. The selection is performed with the help of some logic (the "Selection logic" element of Figure 1).In fact, when stability is detected by more than one phase detection mechanism, more than one bound is estimated and can be selected by the framework.

In this experimental implementation we defined three bound estimators: the conservative estimator, which is always required, and estimators for the exponential and for the Pareto distributions. They are presented in Table 2.

Table 2: Bound estimators for a required coverage.

Estimator	**Minimum time bound t**
Conservative	$bound = E(D) + \sqrt{\frac{V(D)}{1-coverage} - V(D)}$
Exponential	$bound = \frac{1}{\lambda} \ln \frac{1}{1-coverage}$
Pareto	$bound = \frac{k}{\sqrt[\alpha]{1-coverage}}$

5. RESULTS AND DISCUSSION

We did a few experiments to test the implemented framework and to observe, in particular, what would be the improvements achievable with a less conservative and adaptive approach, in comparison with the simple conservative and pessimistic approach described in [2]. This section briefly presents some of our experimental results. The complete test set can be found in an extended version of this paper, available as a technical report [1].

The experiments described here were performed using synthetic data traces with 1500 sample points following the exponential distribution with rate $\lambda = 0.8$. We defined an expected coverage of 0.98 (i.e., given the assumed bound as provided by the framework, no more than 30 timing failures can occur for the 1500 samples of this execution), and a history size of $h = 30$.

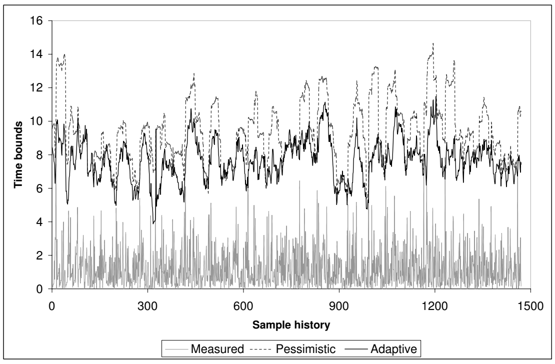

Figure 2: Comparison of pessimistic and adaptive approaches, using "Mechanism 1" for an exponential trace.

A first experiment has been performed aiming to observe the effectiveness of "Mechanism 1" in terms of its capability to detect stable phases, hence providing improved (lower) bounds, still maintaining the required coverage. The results presented in Figure 2 show that this mechanism detected the exponential distribution in 97% of the sample points (we exclude the initial h samples, while the history buffer is not yet full). The average of time bounds produced by the adaptive approach was hence smaller than in the conservative approach (7.6ms against 9.3ms, that is, 18% improvement). In spite of that, we only observed five timing faults in the entire observation period, therefore clearly securing the required coverage.

We note that the adaptive approach makes a pessimistic estimation of $E(\hat{D})$, since it uses the greatest value in the confidence interval of the expected distribution (see [1]). Since the number of observed timing faults was far below the acceptable value, it seems possible to apply a less pessimistic estimation in order to further improve the average achievable time bounds while still ensuring the required coverage.

We executed similar experiments using the same exponential data trace as input, to evaluate "Mechanism 2", "Mechanism 3", and a combination of "Mechanism 1" and "Mechanism 3". Table 3 presents the results expressed in terms of i) percentage of sampling points in which the adaptation mechanism actually detects a stable phase, ii) number of observed timing faults, and iii) average time bounds produced by the framework.

The results show that the three mechanisms present quite similar results when individually applied. However, as expected, a more conservative result is obtained when combining "Mechanism 1" and "Mechanism 3": whenever "Mechanism 1" detects an exponential distribution, the Kolmogorov-Smirnov test is executed to confirm it. This double-checking implementation decreases the aggressiveness of the adaptive approach, thus reducing the ability to detect stability, but still providing a lower average bound (8.1ms) than the conservative approach (9.3ms).

Table 3: Observed results for different mechanisms using the exponential trace.

Mechanism	*Stable detec.*	*Timing faults*	*Avg. Bound*
1	97%	5	7.6 ms
2	96%	5	7.7 ms
3	95%	4	7.8 ms
1 + 3	75%	4	8.1 ms

6. CONCLUSIONS AND FUTURE WORK

In this paper we addressed the problem of supporting adaptive systems and applications in probabilistic environments, from a dependability perspective: maintaining correctness of system properties after adaptation.

We advanced on previous work by leveraging on the assumption that a system alternates stable periods, during which the environment characteristics are fixed, and unstable periods, in which a variation of the environment conditions occurs, *and* that the mode changes can be detected. Based on that, we proposed and evaluated a general framework for adaptation, which allows to dynamically set optimistic time bounds when a stable phase is detected, while it provides conservative but still dependable bounds during transient phases.

We believe that the fundamental conclusion to derive from the full set of experiments made so far, part of them presented in this paper, is that it is possible to define simple and effective mechanisms to detect stable and transient phases and, for the stable ones, correctly characterize the observed probabilistic distribution. Because of that, the proposed framework seems to constitute a promising approach to achieve dependable adaptation and, at the same time, obtain improved (tighter) time bounds than those previously obtained with a more conservative approach. This improvement is relevant in the implementation of practical systems, for instance in the configuration of timeouts in failure detectors, where the objective is to use the smallest possible time bound (to improve the detection time) without compromising the failure detector accuracy (mistakes due to timing faults).

Under the light of the promising findings of this paper, the open question to be addressed next concerns the behavior of the proposed framework under real scenarios. As future work, we plan to validate this approach in a real scenario. We should show that: a) the assumptions stated in this paper are met in real systems; b) the framework is actually configurable for specific environments. We believe this is a hard problem, and thus a very challenging next step.

We are currently extending the framework by incorporating other probabilistic distributions and detection mechanisms, and we plan to apply this framework in the context of a FP6 European project [7], for selected applications involving car-to-car communication in ad hoc environments.

Acknowledgements

We wish to acknowledge the contribution of Odorico Mendizabal, who participated in an earlier version of the document, and to Hans-Peter Schwefel and Hans Reiser, who read earlier manuscripts and provided helpful comments.

7. REFERENCES

[1] A. Casimiro, P. Lollini, M. Dixit, A. Bondavalli, and P. Veríssimo. A framework for dependable adaptation in probabilistic environments. DI/FCUL TR 07–23, Department of Computer Science, University of Lisboa, Oct. 2007.

[2] A. Casimiro and P. Veríssimo. Using the Timely Computing Base for Dependable QoS Adaptation. In *Proceedings of the 20th IEEE Symposium on Reliable Distributed Systems*, pages 208–217, New Orleans, USA, Oct. 2001.

[3] I. M. Chakravarti, R. Laha, and J. Roy. Long-range dependence in variable-bit-rate video traffic. *Handbook of Methods of Applied Statistics*, 1:392–394, 1967.

[4] A. Corlett, D. Pullin, and S. Sargood. Statistics of one-way internet packet delays, Mar. 2002. http://www.potaroo.net/ietf/all-ids/draft-corlett-statistics-of-packet-delays-00.txt.

[5] T. Elteto and S. Molnar. On the distribution of round-trip delays in tcp/ip networks. In *Local Computer Networks, 1999. LCN '99. Conference on*, pages 172–181, 1999.

[6] J. A. Hernández and I. W. Phillips. Weibull mixture model to characterise end-to-end internet delay at coarse time-scales. *IEE Proc. Communications*, 153(2):295–304, Apr. 2006.

[7] HIDENETS Consortium. IST-FP6-STREP-26979 (HIDENETS: HIghly DEpendable IP-based NETworks and Services). Web page: http://www.hidenets.aau.dk.

[8] A. Markopoulou, F. A. Tobagi, and M. J. Karam. Loss and delay measurements of internet backbones. *Computer Communications*, 29(10):1590–1604, June 2006.

[9] M. Menth, J. Milbrandt, and J. Junker. Time-exponentially weighted moving histograms (TEWMH) for application in adaptive systems. In *Proceedings of the Global Telecommunications Conference (GLOBECOM '06)*, pages 1–6, Nov. 2006.

[10] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, and C. Diot. Measurement and analysis of single-hop delay on an ip backbone network. *IEEE Journal on Selected Areas in Communications. Special Issue on Internet and WWW Measurement, Mapping, and Modeling*, 21(6):908–921, Aug. 2003.

[11] N. Piratla, A. Jayasumana, and H. Smith. Overcoming the effects of correlation in packet delay measurements using inter-packet gaps. In *Proceedings of the 12th IEEE International Conference on Networks (ICON 2004)*, pages 233–238, Nov. 2004.

[12] K. S. Trivedi. Probability and statistics with reliability, queuing and computer science applications. John Wiley and Sons, 2002.

[13] P. Veríssimo and A. Casimiro. The Timely Computing Base model and architecture. *Transactions on Computers - Special Issue on Asynchronous Real-Time Systems*, 51(8):916–930, Aug. 2002.

[14] M. Yang, X. R. Li, H. Chen, and N. S. V. Rao. Predicting internet end-to-end delay: an overview. In *Proc. of the 36th IEEE Southeastern Symposium on Systems Theory*, pages 210–214, Mar. 2004.