

On Availability QoS for Replicated Multimedia Service and Content

Giwon On, Jens Schmitt, Ralf Steinmetz

Multimedia Communications (KOM), Darmstadt University of Technology, Germany

Tel.: +49-6151-166150, Fax: +49-6151-166152

Mail: {Giwon.On, Jens.Schmitt, Ralf.Steinmetz}@KOM.tu-darmstadt.de

Abstract. Recently, it has been realized that the importance of satisfying service availability is becoming one of the most critical factors for the success of Internet-based services and applications. In this paper, we take an availability-centric view on QoS where the availability is treated as a new controllable QoS parameter and focus on the issues of providing availability guarantees for distributed and replicated multimedia services and contents. We especially tackle the replica placement (RP) problem and study the effects of number and location of replicas on the achieved availability. From a simulation study, we find that (1) the location of replica is a more relevant factor than their number for satisfying the availability QoS requirements for all individual users, and (2) the heuristic methods, in general, cannot give any guarantee for their achieved availability QoS, while they are very efficient for large size graphs.

1 Introduction

1.1 Satisfying Availability - the Key for Successful Services in Internet

Even though there are many significant research results, technology advances and solutions in *quality of service* (QoS) since the last 20 years [1,2], their application to commercial products or systems was not so successful, in comparison to their attention in the research arena. One critical reason probably is that, as H. Schulzrinne pointed out in [3] and an interview statistic mentioned in [4], the main research focus for QoS was to control transmission characteristics like bandwidth, delay, and loss. This is because applications on the Internet typically assumed in need for QoS support, such as video-on-demand (VoD), tele-teaching, Internet telephony, strongly motivated the development of QoS technologies. While for these the control of the transmission characteristics is certainly important it seems likely by now that, on the one hand, for them this may not be the most pressing need with regard to QoS requirements, and on the other hand that there are other applications having quite different requirements. Indeed, the perceived QoS may be much more influenced by how available a certain service and its data are. In the context of QoS, the availability issue has so far seldom been mentioned, and there is no work known to us which tries to treat availability as a controllable QoS parameter.

1.2 Differentiation of Service Classes and Availability Requirements

While most research efforts in high availability and fault-tolerant systems areas have their focus on achieving the so-called 'five nines' (99.999%) availability [5], there is a demand for service differentiation from service consumers and providers due to costs and competitive nature of the marketplace, which derives from the mechanisms that support different levels of services and their availability. In fact, the need for service differentiation can be observed from different services, such as VoD or news-on-demand over the Internet, as well as different users' requirements which depend on

the service type they demand, on the service time when they access, on the peripherals they have, and on the service price they pay.

From the service system provider's point of view, on the other side, not all system components need to have the same redundancy level, i.e., availability level to offer. The availability level required for individual system components depends on the fact of how much they should be reliable and are critical for offering the service. For example, in developing replication mechanisms for increasing availability of services and their data in a distributed multimedia system *medianode* [6], we analysed the characteristics of multimedia contents and their meta-data and could identify that not all service operations and not all data access functions require the same availability level of the 'five nines'.

1.3 The Main Focus and Approach

The work in this paper is strongly motivated by the two aspects mentioned above. So, the main focus is building a model and mechanisms to study the problem of how to satisfy and guarantee different availability requirements for distributed and replicated multimedia services in a wide-area internetwork like the Internet, and to evaluate the achieved availability QoS. In many existing works, it has been shown that the availability of distributed services and their data can be significantly increased by replicating them on multiple systems connected with each other, even in the face of system and network failures. Thus, we especially tackle the replica placement problem (RPP) and study the effects of number and location of replicas on the reached availability QoS. For this purpose, we develop a concept called *quality of availability* (QoA) in which the availability is treated as a new controllable QoS parameter. Based on the QoA concept, we model the distributed multimedia system as a stochastic graph where all nodes and edges elements are parameterized, statistically independently of each another, with known availability and failure probabilities. We decompose the RPP in three questions: (1) finding a "good" placement for fixed number of replicas, (2) checking the reached QoA with a selected replica placement, and (3) determining the number and location of replicas for satisfying a required QoA with absolute guarantees. Thus, the main focus of the paper is not developing additional, new algorithms for the RPP, but instead specifying the QoA concept and model. For each RP question, we review some solution algorithms from heuristic and exact methods and evaluate their achieved QoA. Based on a simulation study, we find that (1) the location of replica is a more relevant factor than its number for satisfying the required QoA, and (2) the heuristic methods do not give any guarantee for their achieved QoA, even though they are very efficient for large size graphs. Note that we do not address the replica selection and update distribution issues in this work. These issues are handled in our previous work [6] where we also give a comprehensive survey on existing solutions for these problem.

1.4 Outline

The rest of this paper is organized as follows. In Section 2, we develop the concept of quality of availability and describe the QoA metrics to be used in this work. Section 3 presents the replica placement problem and details the algorithms that we reviewed and modified for our problem. In Section 4, we present our implementation methods including the experimental simulation environment and in Section 5 we evaluate the results. Section 6 discusses related work. Finally, Section 7 concludes the paper.

2 The Concept of Quality of Availability (QoA) and its Metrics

2.1 Basic Idea

The basic idea of the QoA concept is that the availability can be defined as *a new controllable, observable QoS parameter*. Indeed, we move the focus of the objective function for the resource and performance optimization problems of the QoS field from satisfying transmission-dependent characteristics such as minimizing transmission delay, jitter, and/or loss to satisfying the availability requirements such as minimizing failure time of service systems and their components and to maximizing the total time amount in which the required service functions as expected and its data are reachable. Given a set of different levels of availability requirements and a network topology with a finite number of possible replica locations, we are then interested in how many replicas are needed, where should they be placed, whether their placement on the given topology satisfies the individually required availability QoS and how they affect the overall service availability quality.

2.2 The QoA Metrics and Parameters

A service is said to be *available* when it functions as expected and its data is reachable. Commonly we distinguish between two levels of available services: *basically available* (BA) and *highly available* (HA). At the BA level, a service delivers correct functionality as long as no faults occur, but it neither offers any redundancy for its system components and data, nor fault detection and recovery mechanisms, while a HA service, in addition to the BA level's feature, provides a certain level of redundancy and eventually the mechanisms for fault-tolerance support [5].

Availability is usually defined either as (a) the percentage of time during which the service is available or (b) the probability of service systems' reachability where each system has an independent failure probability [5]. We use these definitions to specify our availability metrics used in both defining QoA requirements and evaluating reached QoA for networked services. Using these availability metrics - the *percentage of successful service time* and the *failure probability* of underlying systems and network connections, QoA guarantees can be specified in various forms like the cases in network QoS [2]:

- *deterministic* - a service (or its data item) is reachable all the time with an availability guarantee of 99.99 percent. This means for a service that the time duration where the service is unreachable should be absolutely no longer than 53 minutes for a year (1 year = 525600 minutes).
- *probabilistic (or stochastic)* - a service availability probability is guaranteed to be at least, e.g., 90 percent of the whole service access requests.

Actually, the exact form of QoA parameter can be specified both by applications and service providing systems. The QoA evaluation conditions that we use for evaluating satisfied QoA in the evaluation part of this work are as follows:

- *reachedQoA(v)* - it indicates for each demanding node how much the availability requirement has been fulfilled by the selected placement R . For example, the required and satisfied availability values are 95% and 94%, respectively. Then, the reachedQoA is 0.99.
- *minQoA* - it is the minimum of the reachedQoA for all demanding nodes with the selected placement R .
- *avgQoA* - it is the average value of the reachedQoA .

- *guaranteedQoA* - it indicates for how many demanding nodes the selected placement R satisfies the QoA requirement.

Parameter	Notation	Definition
reachedQoA(v)	$QoA_{rch}(v)$	the ratio of satisfied availability to required availability for node v , $\forall v \in V_R$ with $V_R = V$ without R
minQoA	QoA_{min}	$\min \{ QoA_{rch}(v) : \forall v \in V_R \}$
avgQoA	QoA_{avg}	$1/n(\sum QoA_{rch}(v))$ with $\forall v \in V_R$ and $n = (V - R)$
guaranteedQoA	QoA_{gua}	the ratio of $ V_{rch} $ to $ V $, where $V_{rch} =$ set of nodes with $QoA_{rch}(v) \geq 1$

3 The Replica Placement Problem

Distributed, networked service systems that consist of storage/server nodes and network connections between them can be modelled as a graph, $G(V,E)$, where V is the set of nodes and E the set of connection links. This graph is *static* if the members and the cardinality of V and E do not change else it is *dynamic*. The graph is said to be *stochastic* when each node and link are parameterized, statistically independently of each other, with known failure or availability probabilities. The following sections introduce three replica placement (RP) sub-problems that are explored in this work. For all three sub-problems, we model the RP problem as a *static and stochastic* graph.

3.1 Finding a “Good” Placement for a Fixed Number of Replicas

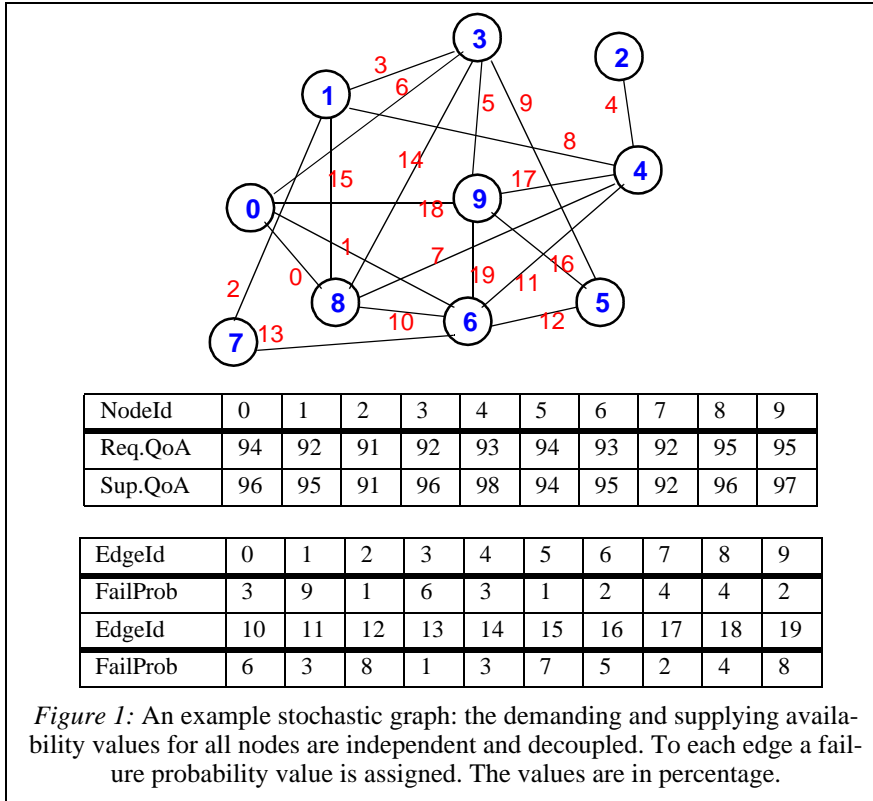
As input, a stochastic graph $G(V, E)$ is given. As the second parameter, a positive integer number k may also be given. The objective of this problem is to place the k replicas on the nodes of V , i.e., find R with $|R| = k$ such that a given optimization condition $O(|R|, R, QoA_condition)$ is satisfied for given availability requirement for service demanding nodes. How well the optimization condition is satisfied depends on the size of $|R|$ and the topological placement R . Because the main goal associated with placing replicas on given networks in our work is satisfying availability QoS which can be required in different levels, we take the availability and failure parameters as our key optimization condition, i.e., $O(|R|, R, guaranteedQoA)$. Thus, with the use of 100% of all clients’, 90%-tile, and mean clients’ required availability value, the optimization condition can be denoted as $O(|R|, R, 1.0)$, $O(|R|, R, .90)$, $O(|R|, R, avgQoA)$, respectively. For these conditions, the replica set R must be chosen such that the maximum, average or any given failure bound for service and its data access meets the QoA requests for any demanding node (client node) of V .

The RP problem can be classified as NP-hard discrete location problem[8,9]. In literature, many similar location problems are introduced and algorithms are proposed to solve the problems in this category. The heuristics such as *Greedy*, *TransitNode*, *Vertex substitution*, etc. are applied to many location problems and have shown their efficiency [10,12,13]. In this work, we take some basic heuristic algorithms as follows. But, different variants of these heuristics and any such improvement can be used with light modifications to enhance the efficiency and performance of our basic heuristics:

- *Random (RA)*. By using a random generator, we pick a node v with uniform probability, but without considering the node’s supplying availability value, and put it

into the replica set. If the node already exists in the replica set, we pick a new node, until the given number reaches k .

- *HighestFirst (HF)*. For each node v , we calculate v 's actual supplying availability value by taking the availability values of all adjacent edges of the node into account. The nodes are then sorted in decreasing order of the actual availability values, and we finally put the best k nodes into the replica set.
- *TransitNode (TR)*. The basic principle of the *TransitNode* heuristic is that nodes with the highest (in/out) degrees, i.e., the number of connection links to adjacent nodes, can potentially reach more nodes with smaller latency. So we place replicas on nodes of V in descending order of (in/out) degree. This is due to the observation that nodes in the core of the Internet that act as transit points will have the highest (in/out) degrees [12].
- *HighestFirst+TransitNode (HF+TR)*: a combination of the *HF* and *TR* algorithms.

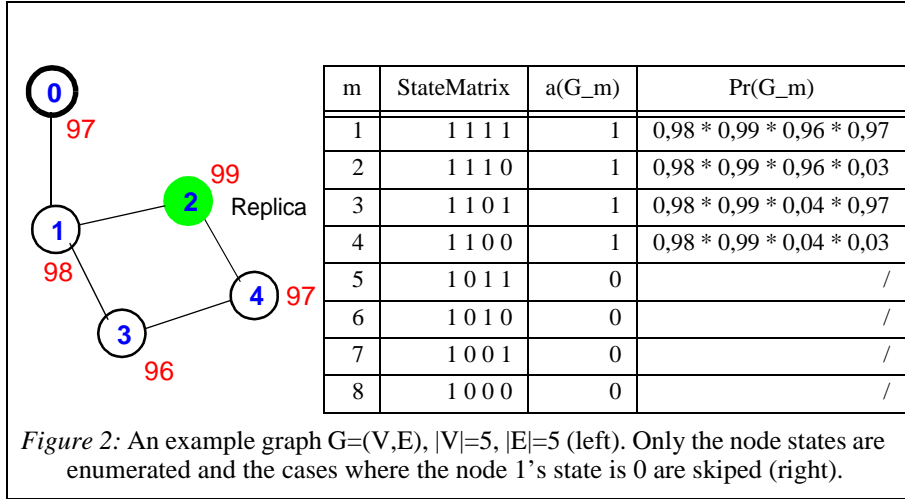


We applied the *HR+TR* algorithm to an example stochastic graph which represents a model of a virtual network topology (Figure 1). In this example graph $G(V,E)$ with $|V|=10$ and $|E|=20$, nodes and links are parameterized with randomly generated probability values, e.g., the demanding and supply availability values of nodes are between 90% and 99%, and the failure probability values of links are between 1% and 10%. For a given number $k = 1$, the candidate replica nodes are $\{4\}$, $\{9\}$, $\{8\}$, $\{3\}$ and $\{0\}$. By considering only the availability value (*Sup.QoA*), the algorithm chooses $\{4\}$ as the replica set.

3.2 Calculating the Reached QoA

Given a stochastic graph $G(V,E)$ and a replica set R . The objective of this problem is to check for all demanding nodes whether the reached availability satisfies the required QoA for them, i.e., whether $\bar{A}(R)$ is 1 or 0. In comparison to the problem of finding a good placement described in Section 3.1, this problem requires a solution which exactly tests whether the result is 1 or 0. Some similar works are introduced in the literature, which are devoted to the problem of *network reliability* [11]. The methods that provide an exact reliability are called exact methods, in contrast to the heuristic methods which provide an approximate result.

Enumerating all possibilities without skipping any solution case requires to take exact methods for solving this problem. From some exact methods which are proposed in the literature, we adopted the state enumeration method [11] and modified it for our problem [7]. In the state enumeration method, the state of each node and each edge are enumerated: the state value is either 1 when it functions or 0 when it fails. Indeed, there are $2^{|V| + |E|}$ states for a graph $G = (V,E)$, i.e., $2^{|V| + |E|}$ partial graphs for G . We then check the QoA for all partial graphs with the replica set given as input.



We applied the *StateEnumeration* algorithm [7] to an example stochastic graph $G(V,E)$ with $|V|=5$ and $|E|=5$, placement $R=\{2\}$, $|R|=1$, as shown in Figure 2. As test node, we take the node 0 which has the availability requirement value 97%. Figure 2 (right) presents the state matrix, the availability and the sum of availability probability value for each partial graph. In this example, we only encounter the states of nodes to reduce the time complexity, i.e., from $2^{|V| + |E|}$ to $2^{|V|}$. The first column means number of the partial graphs to be tested. Instead of considering all partial graph cases of $2^{|V| + |E|}$, we only have the half size of them by skipping the cases in which the node 1 is 'not available', i.e., (the state value of the node 1 is 0), because the node state 'zero' of the node 1 causes no further possible connection for the test node 0 to build any path to the replica node 2. After building a state matrix for the nodes 1,2,3 and 4 at the second column, we check whether there is any path between the node 0 (test node) and the node 2 (replica node) at the third column. According to the result of this check, we calculate the availability probability values for each partial graph of which the availability check value is 1: as shown in the fourth column, we calculated the

$Pr(G_m)$ just for the first 4 partial graphs ($m = 1, 2, 3$ and 4). The summation of the availability probability values of the four satisfied states is: $A_{node0}(G) = Pr(G_1) + Pr(G_2) + Pr(G_3) + Pr(G_4) = 0,97135624$, and this availability value is greater than the availability value required by the node 0. Indeed, the QoA for the node 0 is: $A_{node0}(\{2\}) = 1$.

3.3 Determining the Number & Location of Replicas for Satisfying Required QoA

As input, only a stochastic graph $G(V, E)$ is given. It has to be determined (a) how many replicas must be deployed and (b) where these replicas should be placed to guarantee the required QoA.

Satisfying a certain, required QoA value with a guarantee means that we have to offer always a replica set which fulfils the given QoA requirements in any case. Heuristics are not proper approaches for solving this problem, because they give not always a solution with QoA guarantee. To solve this problem, we generally can use the exact methods, one possible case may be enumeration method such as *StateEnumeration* algorithm which is described in the [7]. For solving this third problem, we need to call the state enumeration method $2^{|V|}$ times, i.e., for all the possible replica solution sets. The algorithm complexity is then $O(2^{|V|} \cdot 2^{|V|} + |E|)$.

4 Simulation Environment

We built an experimental environment to perform a simulation study for the three RP sub-problems addressed in Section 3. Our goal in conducting an availability evaluation is to determine exactly the replica set R for given QoA requirements, and to study the effect of changing R and $|R|$ on the required and reached QoA which is given as the optimization condition $O(|R|, R, a)$, on the other side. For our availability evaluation, we conducted simulations on random network topologies.

By using Leda graphic library [14] several random topologies in different sizes can be generated at run time. We also used graph files which are generated by the topology generator Tiers[15]. To reflect the actual Internet topology, we used several different network trace data from NLNR[16], which describe the Autonomous System (AS) topology on a different day, and associated in/out-degree of the AS with each node of the randomly generated graphs. But, because we had no access to any real data concerning the availability or failure probability parameters, we simply assign the values from a certain range to each node and each edge.

The availability and failure probability parameters for nodes and edges of the graphs are single values: for example, 50, 80, 90 or 99% as availability values and 10, 5, or 1% failure probability values. We decoupled the availability values between the demanding and supply nodes, i.e., all nodes have two availability parameters assigned: one value as the demanding availability parameter and the other as the supplying. Thus, when a node is a demanding node (client node), then its demanding availability value is used, while for supplying node the actual supplying availability value is, for example, calculated by multiplying the availability values of its supplying own and the average value of its adjacent edges. At the replica set building phase, each node is evaluated according to its supplying availability value. Thus, to be elected as a replica node, for example in the *HA* algorithm, a node should have a high supplying availability value.

As replication model, we assume the *full replication* in which the whole data items of an origin server system are replicated to other nodes located within the same network.

Mirroring is a typical case of the full replication model. The simulation program is written in C/C++ and runs on Linux and Sun Solaris 2.6 machines.

5 Evaluation of Reached QoA

In this section, we present our experiment results. We evaluated the reached QoA of our heuristics and the exact enumeration method using topologies of different sizes. We ran each basic heuristic and the exact state enumeration method on each topology using different value ranges for the availability and failure probability parameters of nodes and edges. The demanding and initial supplying availability values of the nodes and the failure probability values of the edges are assigned randomly, from a uniform distribution where we varied the parameter values as Table 1 shows. To evaluate the QoA offered by our heuristics and *StateEnumeration* algorithm, we used the QoA metrics defined in Section 2.2.

Table 1: QoA parameters and their values for our simulaiton

Type	Parameter	Notation	Value
Graph	node and edge size	$ V : E $	G1(20:30), G2(100:300)
Edge	edges' failure probability	q_l	1 ~ 10 %, 0%
Node	nodes' required availability	P_{nReq}	90~99%, 50~99%, 50-90-99%
	nodes' supply availability	P_{nSup}	$P_{nSup} \geq P_{nReq}$

5.1 Relative Comparison of Reached QoA by Heuristics

We evaluate at first the reached QoA by our simple heuristics. The baseline for our experiment is an initial placement R_0 which is obtained by randomly selecting k nodes from V . We then compare the reached QoA of each heuristic to this baseline and present the relative QoA improvement obtained with each heuristic.

5.1.1 Effects of R and |R| on Reached QoA

The first experiment was to find good locations of a replica set R with $|R| = k$ for given graphs G with maximal replica number k . The conditions that we assumed for this problem were: (1) $QoA_{min} > 0.9, 0.95, \text{ and } 0.99$, respectively, and (2) $QoA_{avg} > 1.0$. In this case, there was no constraints on the topological location of the replicas and replicas may be placed at any node v in G .

Figure 3 (*top*) shows the results from this experiment with $G2$. We plot the number of k on the x-axis and the reached QoA on the y-axis. In each graph, we plot different curves for different heuristics and different ranges for required availability values. From Figure 3 (*top*), we can see that our heuristics HA and $HA+TR$, although they are very simple, reach significantly higher QoA in comparison to the baseline placement. Even though the improvement of 12% QoA guarantee rate with replicas 5 to 25 (totally, 20% of the whole nodes are replicas) may not seem much, it is important to note that the number of replicas is really a relevant factor for improving QoA: the larger the replica number is, the better is the reached QoA.

5.1.2 Effects of Varying Availability Requirement Value Ranges on QoA

In the second experiment, to study the effects of different ranges of the required availability values on the reached QoA, we varied the ranges of required availability values (P_{nReq}) from 50~99% to 80~99% and 90-99% for the same graph $G2$. We took

also P_{nReq} as different single values like 50-80-99%. As Figure 3 (bottom) shows, the improvement rate of the reached QoA is better when the P_{nReq} is distributed in wider ranges.

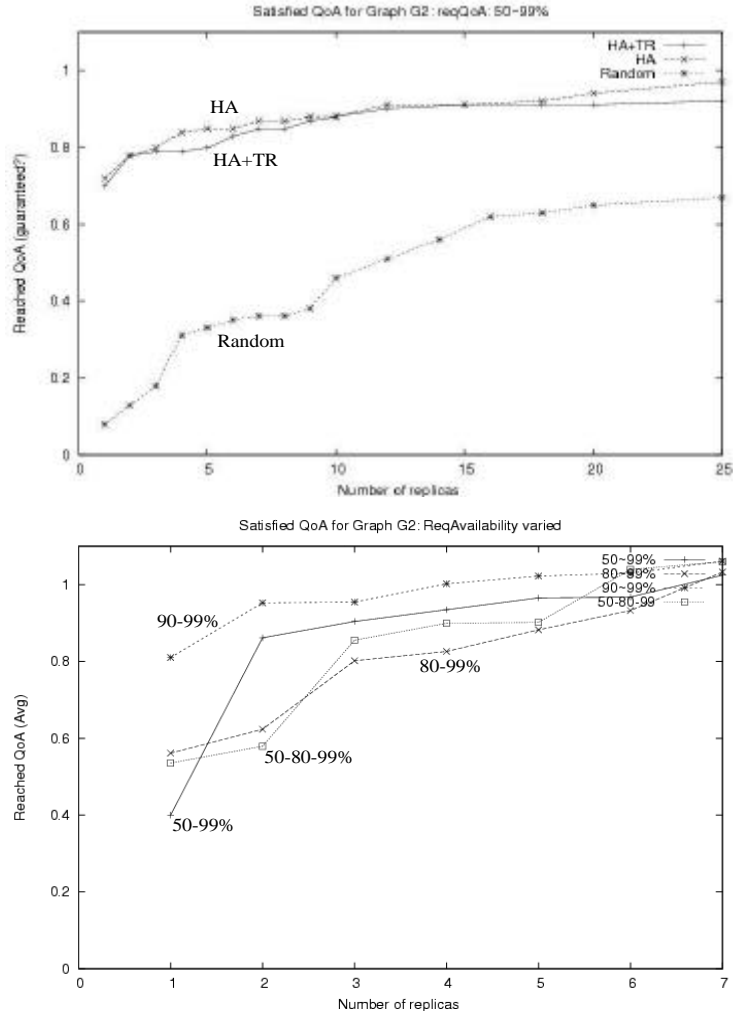


Figure 3: Reached QoA values by our heuristics: (top) different heuristics, (bottom) different ranges for required availability values

5.2 Exact Evaluation of Reached QoA by Heuristics

We now evaluate the QoA reached by our heuristics in an exact form and check whether the reached QoA can really satisfy the required QoA for all demanding nodes (*test for guaranteedQoA*). We test also how many replica nodes do the heuristics need to give a QoA guarantee, i.e., *guaranteedQoA* = 1. For this purpose we ran our *StateEnumeration* routine with replica sets produced by our heuristics *HA* and *HA+TR* as input. Due to the exponentially growing runtime complexity and the memory

requirements with growing graph sizes, we limited our experiments for the *StateEnumeration* to a small graph, the test graph *GI* with $|V| = 20$ and $|E| = 30$. Table 2 shows the detailed test result from *HA+TR*. For the calculation of the average (*avgQoA*) and minimal reached QoA (*minQoA*), we excluded the QoA values for replica nodes.

Table 2: A detailed result for *HA+TR* with *GI*, $q_l : 0\%$, and P_{nReq} range: 90-99%

$ R/$	Replica locations, R	Avg QoA	Min QoA	QoA_{gua}	QoA_{gua} by exact test
1	8	1.0118	0.9100	0.75	0.80
2	8,10	1.0194	0.9100	0.75	not checked
3	8,10,12	1.0226	0.9193	0.75	not checked
4	8,10,12,11	1.0355	0.9496	0.85	not checked
5	8,10,12,11,13	1.0399	0.9496	0.85	not checked
6	8,10,12,11,13,0	1.0487	0.9591	0.90	0.95
7	8,10,12,11,13,0,16	1.0556	0.9900	0.95	1.00
8	8,10,12,11,13,0,16,1	1.0577	0.9900	0.95	not checked
9	8,10,12,11,13,0,16,1,2	1.0610	0.9900	0.95	not checked
10	8,10,12,11,13,0,16,1,2,5	1.0711	1.0000	1.00	1.00

Even though *HA+TR* could reach the average QoA (1.0118) greater than 1 with one replica node, it could not offer the QoA guarantee: 10 replicas were needed to satisfy the QoA guarantee for the small graph.

5.3 Finding the Optimum - $|R/$ and R

In the last experiment, we considered the case of finding the optimum, i.e., the minimal number of replicas and their geographical placement which satisfies the availability QoS with guarantee. We re-used *StateEnumeration* and the test graph *GI* with the same values for availability and failure probability parameters. We started the routine with a replica degree of 1, i.e., $k=|R/ = 1$, and selected each node as replica node. We then incremented the replica degree, until we reached the $QoA_{gua} = 1.0$ (a QoA with guarantee). Table 3 shows the reached QoA values at each k ($k=1,2,3$). Figure 4 plots the reached QoA that *StateEnumeration* algorithm calculated exactly with each instance for the given k . The wider spectrum of the left part is for highlighting the reached QoA from all of the instances for $k=1$ and the chosen instances for $k=3$. The right part of Figure 4 shows how the reached QoA varies in the case of $k=2$, and how big is the gap between good and bad QoA rates reached by the instances.

5.4 Discussion

The following observations could be identified from our experiment results: (1) the location of replicas is a relevant factor for the availability QoS. Even though the QoA improvement could be achieved by increasing replica numbers, replicas' placement and their dependability affected the QoA more significantly; (2) using a heuristic method is more efficient than the exact method, at least in terms of the runtime complexity, to find a good placement for large graphs. But, the replica degree of their placement results are in most cases higher than those of exact methods. Furthermore,

the heuristics give no guarantee for availability QoA; (3) in opposite to the heuristic method, the exact method guarantees the availability QoS with its placement results, although the runtime complexity is very high: $O(|V_R| \cdot 2^{|V|+|E|})$ and $O(2^{|V|} \cdot |V_R| \cdot 2^{|V|+|E|})$ for the availability checking (Section 3.2) and the guaranteed QoA problems (Section 3.3), respectively.




Figure 4: Reached QoA that was checked exactly by *StateEnumeration*

Table 3: A test result from *StateEnumeration* with *G1*, failure probability (q_l): 0%, and req. availability (p_{nReq}) range: 90-99%

No. of replicas	Best QoA value	Worst QoA value	Mean QoA value	Instances achieved the best QoA value
1	0.80	0.10	0.3345	{0},{8}
2	0.95	0.15	0.8078	{0,11},{0,18},{8,11},{8,18},{11,13}, {12,16},{13,16}
3	1.00	not checked	not checked	{0,11,16},{0,16,18},{8,11,16}, {8,16,18},{11,12,16},{11,13,16}

6 Related Work

The key ideas on which our work on QoA concept in this paper bases are (i) an availability-centric view on QoS and (ii) satisfying different levels of QoA values required by individual users. Since the common goals associated with placement problems in existing studies are reducing clients' download time and alleviating server load, the main feature of the problem solving approaches for this problem category is that they usually addressed the cost and resource minimization issues, but not the question how to guarantee the required availability. Furthermore, we can find an "good" upper bound, if the selected placement meets the required availability QoS, but it is not guaranteed that the selected placement always meets the availability requirement.

7 Conclusion

In this paper we took an availability-centric view on QoS and focused on the issues of providing models and mechanisms to satisfy availability requirement for replicated multimedia services and contents. We developed a concept called *quality of availability (QoA)* in which the availability is treated as a new controllerable QoS parameter. Based on the QoA concept, we modelled a distributed multimedia system as a stochastic graph where all nodes and edges are parameterized with known availability and failure probabilities. We especially tackled the replica placement problem in which we specified different placement problems with different QoA metrics such as *minQoA*, *avgQoA*, and *guaranteedQoA*. The primary result shows already that (1) the location of replicas is a more relevant factor than their number for satisfying the availability QoS for different users, and (2) the heuristic methods could not give any guarantee for their achieved QoA, even though they are very efficient for large size graphs. Our proposed QoA concept and model can be used as a base mechanism for further study on the availability and reliability QoS with dynamic replication problems and mobile storage planning problems.

References

- [1] Zheng Wang. *Internet QoS: Architectures and Mechanisms for Quality of Service*. Lucent Technologies, 2001.
- [2] J. Schmitt. *Heterogeneous Network QoS Systems*. PhD thesis, Darmstadt University of Technology, December 2000.
- [3] H. Schulzrinne. "QoS over 20 Years". Invited Talk in *IWQoS'01*. Germany, 2001.
- [4] SEQUIN project. "Quality of Service Definition" *SEQUIN Deliverable D2.1*, April 2001. Available at <<http://www.dante.net/sequin/QoS-def-Apr01.pdf>>
- [5] G. Coulouris, J. Dollimore, and T. Kindberg. *Distributed Systems*. 3rd Ed., Chapter 1, 8, 14 and 15, Addison-Wesley, 2001.
- [6] G. On, J. Schmitt and R. Steinmetz. "Design and Implementation of a QoS-aware Replication Mechanism for a Distributed Multimedia System", in *Lecture Notes in Computer Science 2158 (IDMS 2001)*, pp.38-49, Sep. 2001.
- [7] G. On, J. Schmitt and R. Steinmetz. "The Quality of Availability: Tackling the Replica Placement Problem", *Technical Report*, KOM-TR-2001-11, Darmstadt University of Technology, December 2001.
- [8] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.
- [9] Christos H. Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, ISBN 0-13-152462-3, 1982.
- [10] N. Mladenovic, M. Labbe and P. Hansen: "Solving the p-Center Problem with Tabu Search and Variable Neighbourhood Search", July 2000.
- [11] C. Lucet, J.-F. Manouvrier. "Exact Methods to compute Network Reliability". In *Proc. of 1st Inter'l Conf. on Math. Methods in Reliability*, Bucharest, Sep. 1997.
- [12] S. Jamin, C. Jin, A. R. Kurc, D. Raz, Y. Shavitt. "Constrained Mirror Placement on the Internet", In *Proc. of IEEE INFOCOM'01*, pp. 31-40, 2001.
- [13] P. Krishnan, D. Raz and Y. Shavitt. "The Cache Location Problem", In *IEEE/ACM Transactions on Networking*, 8(5), pp. 568-582, Oct. 2000.
- [14] Leda. software available from <<http://www.algorithmic-solutions.com/>>
- [15] Tiers. software available at <<http://www.isi.edu/nsnam/dist/topogen/>>
- [16] NLANR measurement and operations analysis team. <<http://moat.nlanr.net>>