Availability of enterprise IT systems – an expert-based Bayesian model

Ulrik Franke, Pontus Johnson, Johan König and Liv Marcks von Würtemberg
Industrial Information and Control Systems
Royal Institute of Technology
Stockholm, Sweden
{ulrikf, pj101, johank, livm}@ics.kth.se

***Abstract*—Ensuring the availability of enterprise IT systems is a challenging task. The factors that can bring systems down are numerous, and their impact on various system architectures is difficult to predict. At the same time, maintaining high availability is crucial in many applications, ranging from control systems in the electric power grid, over electronic trading systems on the stock market to specialized command and control systems for military and civilian purposes.**

The present paper describes a Bayesian decision support model, designed to help enterprise IT systems decision makers evaluate the consequences of their decisions by analyzing various scenarios. The model is based on expert elicitation from 50 academic experts on IT systems availability, obtained through an electronic survey.

The Bayesian model uses a leaky Noisy-OR method to weigh together the expert opinions on 16 factors affecting systems availability. Using this model, the effect of changes to a system can be estimated beforehand, providing decision support for improvement of enterprise IT systems availability.

***Keywords*-Systems availability, High availability, Downtime, Bayesian networks, Noisy-OR, Expert elicitation**

I. INTRODUCTION

Maintaining high enterprise IT systems availability is a high priority throughout many industries. In a frequently cited report from 1998, IBM Global Services report that unavailable systems cost American businesses $4.54 billion in 1996, due to lost productivity and revenues [1]. The report goes on to list average costs per hour of downtime ranging from airline reservations at $89.5 thousand to brokerage operations at $6.5 million (all in 1998 dollars). A vivid reminder of the financial sector's sensitivity occurred when the Nordic and Baltic stock markets were forced to close down for 5 hours on June 4 2008 due to the trading system Saxess being down. This outage prevented transactions worth approximately 20 billion SEK (ca €2 billion) [2].

While useful in some contexts, cost estimates do not always accurately reflect the criticality of systems availability. This is often the case for IT systems supporting emergency response, police and military operations, etc. In a recent Gartner report, it is therefore recommended that investments to ensure high availability in such systems are justified using qualitative measures of the impact on the population affected [3]. The same line of reasoning applies to information and control systems serving critical infrastructure, such as the electric power grid, railway transportation, water supply etc.

Citing the data on average downtime costs referred to above, Marcus and Stern [4] observe that not all losses are easy to quantify, in particular when they are partly composed of opportunity costs, as in the case of brokerage services. However, this is not to denigrate the importance of availability, as they go on to list some indirect costs that can be brought about by system outages: (i) poor customer satisfaction, (ii) bad publicity, (iii) plummeting stock price (while [5] suggests that this effect is actually small), (iv) legal liabilities, (v) worsened employee morale and (vi) an impact to external reputation.

Yet another measure on the importance of systems availability is stakeholder polling. In a recent survey, 178 enterprise IT system executives and practitioners from Sweden and the German-speaking countries were asked to assess future prioritization of various system qualities (in the sense of ISO-9126 [6]) in their companies. On a five point Likert scale, 48.9% of respondents gave availability the highest mark, making it the most highly prioritized system quality in the survey [7]. A Gartner report, based on surveys conducted in 2007 and 2008, notes consistent findings and concludes that "[t]he overall proportion of mission-critical IT services continues to increase, along with the cost of business downtime" [8].

A. Outline

The remainder of the paper is structured as follows: Section II contrasts the present contribution with some related work. Section III introduces Bayesian networks, survey methodology in general, and the particular methods employed for learning Bayesian networks from experts. Section IV is the locus of the main contribution. Here, the results of the expert survey are described, and the resulting Bayesian model for assessment of Enterprise IT systems availability is built. An example of its usage is presented in section V. A discussion of the strengths and weaknesses of the contribution then ensues in section VI, followed by some concluding remarks in section VII.

II. Related work

A general and widely cited description of IT systems availability is found in [4], where the authors present an "availability index" describing the relationship between various availability-increasing measures and their costs. The presented availability index gives guidance on improving systems availability, but it is not empirically validated in a structured way. The present contribution partially aims to address this by taking [4] as the basis for the survey questions, as discussed in section III-B.

In [9] the authors present an approach for analytical service availability assessment, mapping dependencies between low-level ICT infrastructure and high-level services. The mapping, however, does not give a detailed description of the supporting ICT infrastructural elements, nor any weighting of how each element impacts the service availability. In [10] a similar mapping is presented, but here the focus is the impact of ICT infrastructure availability upon business processes, rather than upon availability assessment as such.

An effort to identify factors impacting software reliability is presented in [11]. The article includes the identification of 32 factors involved in the software development process, all of which impact software reliability. A ranking based on empirical research from 13 companies working with software development is presented, highlighting the most important factors influencing the software reliability. However, only the software development phase is addressed – how to ensure availability once systems already taken into service is not mentioned.

The application of Bayesian networks for information system quality analysis is proposed and applied in [12]. In this paper, an enterprise architecture evaluation framework for the analysis of information systems modifiability is presented. An expert survey was conducted in order to create a Bayesian model, the details of which are found in [13]. The present paper is similar in method, but focuses on availability rather than modifiability.

III. Method

A. Bayesian networks

Friedman et al. [14] describes a Bayesian network, $B = (G, P)$, as a representation of a joint probability distribution. The first component G is a directed acyclic graph consisting of vertices, V, and edges, E, i.e. $G = (V, E)$. The vertices denote a domain of random variables $X_1, \dots, X_n$, also called chance nodes. Each chance node, X_i, may assume a value x_i from the finite domain $Val(X_i)$. The edges denote causal dependencies between the nodes, i.e. the causal relations between the nodes. Whenever an edge goes from a node X_i to a node X_j, X_i is said to be a causal parent of X_j. The second component P of the network B, describes a conditional probability distribution for each chance node, $P(X_i)$, given the set of its causal parents $\text{Pa}(X_i)$ in G. It is now possible to write the joint probability distribution of the domain $X_1, \dots, X_n$ using the chain rule of probability, in the product form:

$$P(X_1, ..., X_n) = \prod_{i=1}^{n} P(X_i|\text{Pa}(X_i)) \tag{1}$$

In order to specify the joint distribution, the respective conditional probabilities that appear in the product form must be defined. The component P describes the distribution for each possible value x_i of X_i, and $\text{pa}(X_i)$ of $\text{Pa}(X_i)$, where $\text{pa}(X_i)$ is the set of values of $\text{Pa}(x_i)$. These conditional probabilities are represented in matrices, here forth called Conditional Probability Distributions (CPDs). Using a Bayesian network, it is possible to answer questions such as: what is the probability of variable X being in state x_1 given that $Y = y_2$ and $Z = z_1$.

In the general case, the relations between variables described by the conditional probability matrices can be arbitrarily complicated conditional probabilities. The model presented in this paper uses only a single rather simple relation, leaky Noisy-OR, described in section III-C.

More comprehensive treatments on Bayesian networks can be found in e.g. Neapolitan [15], Jensen [16], Shachter [17] and Pearl [18].

B. Expert elicitation

Expert elicitation is the process where a person's knowledge and beliefs about one or more uncertain quantities are formulated into a joint probability distribution [19], i.e. the act of parameter estimation through the use of domain experts. This approach is generally used when available datasets are sparse in comparison with the number of nodes that need to be parameterized [20]. Using a well-structured process for expert elicitation is important in order to minimize the bias of the domain expert. A rough outline of such an elicitation process is given in [21]:

1) Select and motivate the expert
2) Train the expert
3) Structure the questions
4) Elicit and document the expert judgments
5) Verify the results

In the following, we detail how each of these steps were carried out in the present study.

1) Select and motivate the expert: The selection of respondents in the present survey was based upon academic publications. To identify respondents, searches were performed in major publishing databases (Springer and Elsevier), in professional societies databases such as the IEEE and in pure indexing databases such as SCOPUS. The search criteria involved combinations of topic-words such as "availability", "reliability" and "dependability" with research area delimitations such as "information system", "IT system" and "corporate IT". The resulting selections of articles were then manually screened, based on title and abstract (if sufficient) or full content (if necessary) to determine whether the authors should be invited to participate or not. Whenever several co-authors to a single paper were encountered, no distinction was made between them (all or none were invited). The searches were limited in time to the past decade, i.e. only publications from 1999 and onward were selected. In all, 154 authors of journal articles, 298 authors of conference articles and 11 authors of edited volumes were invited to participate, i.e. a grand total of 463 experts.

As the experts consulted in this study were widely geographically spread, a mail survey was used [22]. Another reason to use a mail survey is that the non respond bias of mail surveys tends to be directly related to the subject, i.e. chiefly respondents particularly interested in the subject return the questionnaire [23]. This effect will be further discussed in section IV. The internet-based application SurveyMonkey hosted the survey, which was open for two weeks, from January 4 to January 15, 2010. As recommended in [24], a reminder was sent to non-responding participants in the middle of the second week to increase the response rate.

As noted in [21], it is important to convince the experts that there is no straightforward way to tell a ”right” from a ”wrong” answer, but that their assessments should only represent their very own knowledge and experience as faithfully as possible. Indeed, the very rationale for selecting this particular research approach is that the subject is difficult to investigate in other ways. In the introduction to the present survey, it was therefore clearly stated that ”your particular piece of experience and your corresponding answers are very important to us as we try to build a general model”. Furthermore, each question in the present survey included a self-assessment on the credibility of the answer, enabling anyone feeling uncertain to communicate this. As will be discussed in section III-C, this self-assessment also plays an important role in the construction of the Bayesian model.

2) Train the expert: The validity of the study is highly dependent on the respondents’ comprehension of the questions. Therefore, it is often advisable to spend part of the survey to train the expert [25], so that she will not only be a subject matter expert, but also an expert in giving probability estimates. In the present survey, this was accomplished by the use of an initial tutorial question, where the scope and aim of the question was explained at some length using text and figures.

During the training phase, feedback on answers with known correct answers can help experts calibrate their responses [26]. However, in the present study, this was not feasible due to lack of indisputable data of sufficient generality.

3) Structure the questions: There are some different approaches to elicitation, direct elicitation being the most obvious and straightforward one. Here, questions are asked along the lines of ”What is the probability that variable A takes this state given these parent values?”. However, these questions can be hard for domain experts to relate to [19], forcing the use of alternative approaches as described for example in [27]. In the present survey, a behaviorally anchored scale [22] was used.The experts were asked to answer ”How large a share of currently unavailable enterprise IT systems would you guess would be available if a best practice *factor X* had been present?” (Mutatis mutandis, depending on the appropriate grammar of each factor.) The factors themselves were derived from the availability index presented in [4]. Their completeness is thoroughly discussed in section VI.

Including subjective formulations such as ”best practice” in a survey has both advantages and disadvantages. On the one hand, it is possible to interpret the question in several ways which makes it more difficult to compare the answers with each other. On the other hand, the answers are to a lower extent limited to assumptions specified in the question [22]. In this particular case, the authors of this article do not claim to know the technical details of ”best practice” better than the respondents, which is the reason why a subjective formulation was used.

A separate question was written for each probability to be assessed. In those cases were potentially ambiguous or unclear terms were used, a short explanatory note was appended to the question to clarify the intended use. To provide the respondent with a birds-eye-view of the survey, a figure illustrating all question categories in order of appearance was continuously displayed throughout the survey.

As noted in [28], experts dislike writing numbers for subjective probabilities and prefer to check scales, place an ’X’ in a box, etc. In the present survey, this was accommodated by using predefined scales in drop-down lists for the alternatives.

To make the survey questions as clear and lucid as possible, a few test surveys were tried out iteratively, as recommended in [25]. The test respondents included both non-domain experts (for general advice on structure and readability) and a two of the actual respondents (for more topic-related advice).

4) Elicit and document the expert judgments: Since elicitation is taxing for the expert, [28] recommends that sessions should not exceed one hour. The present survey being web-based, with the possibility for the respondent to take a break or withdraw at his or her discretion, this problem can be considered of marginal importance. However, if a survey is too long or too complex, the response rate of the questionnaire decreases [24]. The level of detail in this study was therefore limited by the expected response rate considered acceptable. A survey of the responses ex post indicates that a typical full response required about 20-30 minutes.

5) Verify the results: This aspect is addressed in section VI below.

C. Building Bayesian networks

Bayesian networks are a powerful formalism, but their use requires the specification of conditional probability distributions (CPDs). As the number of variables $X_1, \ldots, X_n$ causally affecting a target variable Y grows, fully specifying these distributions becomes increasingly cumbersome. As noted by [29], a binary variable with n causal parents requires 2^n independent parameters to exhaustively describe the conditional probabilities. As n grows, 2^n parameters quickly becomes a prohibitive number. Often, however, canonical parameter-based distributions can be used to decrease the modeling effort, yet still give a sufficiently good approximation of the true distribution [29].

The solution described in [29] is the use of a Noisy-OR gate. Using this formalism, the number of parameters required from expert estimation becomes only n, a significant gain. The underlying assumption is that instead of investigating every combinatorial interaction among the $X_1 \ldots X_n$ causal parent variables, their interactions are modeled by a Noisy-OR gate. Furthermore, since Noisy-OR distributions approximate CPDs using fewer parameters, the resulting distributions are in general more reliable, being less susceptible to overfitting [30].

The Noisy-OR gate [18], [29] is typically used to describe the interaction of n causes $X_1, \ldots, X_n$ to an effect Y. In the present article, of course, this effect Y is the unavailability of enterprise IT systems. Two assumptions are made, viz. (i) that each of the causes has a probability p_i of being sufficient for producing Y, and (ii) that the ability of each cause X_i, to bring about Y is independent. Mathematically, the following holds:

$$p_i = P(y|\bar{x}_1, \bar{x}_2, \ldots, x_i, \ldots, \bar{x}_n) \tag{2}$$

where x_i designates that causal factor X_i is present and $\bar{x}_i$ that it is absent. It follows that the probability of y given that a subset $\mathbf{X}_p \subseteq \{X_1, \ldots, X_n\}$ of antecedent causes are present can be expressed as:

$$P(y|\mathbf{X}_p) = 1 - \prod_{i:X_i \in \mathbf{X}_p} (1 - p_i) \tag{3}$$

This is a compact specification of the CPD.

A natural extension proposed by Henrion [31] is the so called *leaky* Noisy-OR gate. The rationale for the leakage is that models typically do not capture all causes of Y. If some potential X_i have been left out, as is the case in "almost all situations encountered in practice" [29], this shortcoming can be reflected by adding an additional parameter p_0, the leak probability, such that

$$p_0 = P(y|\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n) \tag{4}$$

In words, this reflects the probability that Y will occur spontaneously, in the absence of all the explicitly modeled causes $X_1, \ldots, X_n$. In a leaky Noisy-OR gate, the CPD is becomes

$$P(y|\mathbf{X}_p) = 1 - (1 - p_0) \prod_{i:X_i \in \mathbf{X}_p} \frac{(1 - p_i)}{(1 - p_0)} \tag{5}$$

The aim of the expert elicitation survey can now be stated more explicitly. For each of the factors $X_1, \ldots, X_n$ identified in the survey, a probability p_i of X_i being a cause of enterprise IT system unavailability can be estimated. Depending on the respondent comments regarding causes not listed in the survey, an approximate value of p_0 can also be found, as discussed in section IV.

IV. Results

A. The respondents

Figure 1 displays some basic data about the respondents who chose to start taking the survey, viz. their affiliation, professional relation to enterprise IT systems availability and their number of years of experience of enterprise IT systems availability.

The data illustrated pertains to two groups: the 96 respondents who began the survey, and the 50 who completed it. For obvious reasons, no similar data exists for the grand total group of the 463 experts invited. As seen in Figure 1, there is no obvious change in the proportion of affiliations from the group that began the survey to the group that completed it. As for working with questions related to enterprise IT systems availability, the reassuring trend is that those not involved in the field to a large extent dropped out from the survey. Thus, while 72% (18 out of 25) of those who identified themeselves as "to a large extent" working with enterprise IT systems availability completed the whole survey a mere 25% (5 out of 20) of those "not at all" working with enterprise IT systems availability did so. Based on these figures, it seems reasonable to assume that the quality of the responses collected was improved by this self-selection. A similar trend can be seen when it comes to the number of years of experience.

B. The Bayesian model

The main results of the survey are presented in Table I. As described in section III, the respondents not only answered each question, but also stated their certainty. In the table, the N specified excludes respondents who "just picked a random interval", and retains only those who were 50%, 90% or 99% certain. This corresponds to the answers actually used in building the Bayesian model.

As can be seen, the number of useful answers vary from 36 to 54 over the causal factors, with a mean of 44. While Table I does give a good overview of the results, it does not show the self-assessed uncertainties associated with each

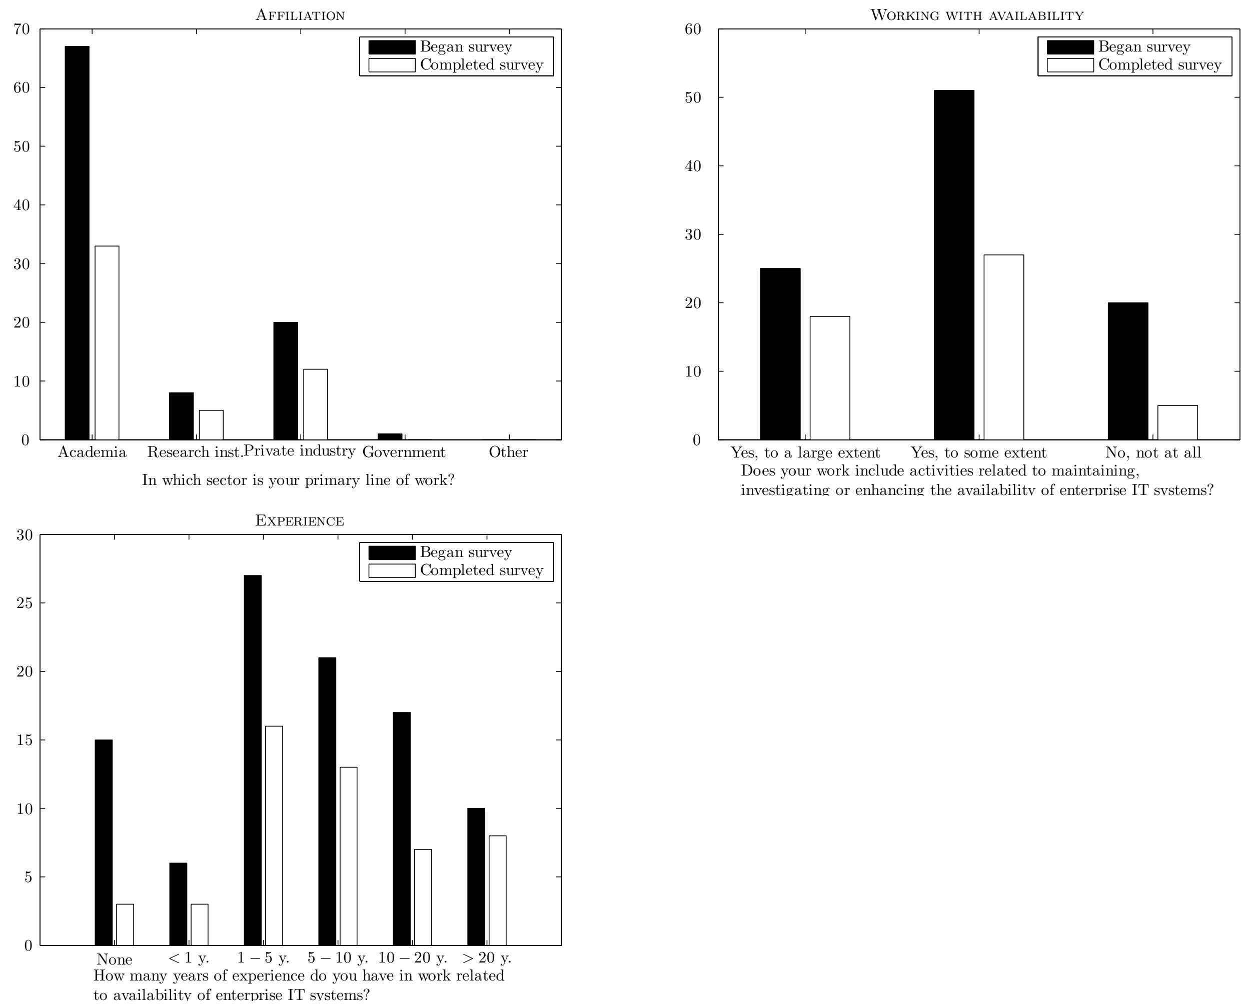

Figure 1. Data on respondents' affiliations, working experience with enterprise IT systems availability, and years of such experience.

Table I
CAUSAL FACTORS (BASED ON [4]) AND STRENGTHS AS PER THE RESPONDENTS' ANSWERS.

		How large a share of currently unavailable enterprise IT systems would you guess would be available if a best practice *factor X* had been present?								
	Causal factor X	$< 0,05\%$	$0,05\%-0,1\%$	$0,1\%-0,5\%$	$0,5\%-1\%$	$1\%-5\%$	$5\%-10\%$	$10\%-50\%$	$> 50\%$	N
1	Physical environment	1	2	3	5	8	18	14	3	54
2	Requirements and procurement	1	1	1	2	10	11	16	5	47
3	Operations	1	0	0	1	6	16	19	5	48
4	Change control	0	1	1	2	7	13	17	8	49
5	Technical solution of backup	2	3	1	6	9	15	5	3	44
6	Process solution of backup	0	2	4	6	12	10	9	0	43
7	Data redundancy	0	2	4	6	8	11	9	4	44
8	Storage architecture redundancy	0	3	3	8	15	6	4	1	40
9	astructure redundancy	1	3	3	8	12	7	5	3	42
10	Avoidance of internal application failures	0	4	3	2	7	12	15	2	45
11	Avoidance of external services that fail	2	1	0	5	4	15	12	4	43
12	Network redundancy	0	4	1	4	11	13	9	2	44
13	Avoidance of network failures	0	5	1	5	9	9	11	2	42
14	Physical location	2	0	7	7	11	8	3	2	40
15	Resilient client/server solutions	1	1	2	5	9	9	6	3	36
16	Monitoring of the relevant components	0	2	1	4	13	6	14	3	43

answer by the respondents. These, however, play a vital role in determining the Noisy-OR probabilities $p_1, \ldots, p_n$ associated with each of the causes $X_1, \ldots, X_n$ listed in the table.

Table II
THE 54 USEFUL RESPONDENT ANSWERS REGARDING THE PHYSICAL ENVIRONMENT FACTOR, DISPLAYED BY CERTAINTY.

Physical environment	50% (I think so.)	90% (I am quite sure.)	99% (I am almost completely certain.)	Weighted vote w
$< 0,05\%$	0	0	1	0.99
$0,05\% - 0,1\%$	2	0	0	1
$0,1\% - 0,5\%$	2	1	0	1.9
$0,5\% - 1\%$	4	0	1	2.99
$1\% - 5\%$	6	2	0	4.8
$5\% - 10\%$	16	2	0	9.8
$10\% - 50\%$	10	3	1	8.69
$> 50\%$	1	2	0	2.3

Table II illustrates the distribution of answers over the intervals with certainty gradings for the first causal factor, physical environment. As before, respondents who "just picked a random interval" are excluded. As can be seen, most feel comfortable with the 50% level, "I think so".

To weight these judgments into a single probability p_i for use in the Noisy-OR model, the number of respondents in an interval j has been multiplied with the certainty $q \in \{0.5, 0.9, 0.99\}$ of their responses (these figures were the alternatives used in the survey). The weighted voting score w_j of interval j is thus defined as

$$w_j = \sum_{k \in K_j} q_k = 0.5 \cdot n_{0.5} + 0.9 \cdot n_{0.9} + 0.99 \cdot n_{0.99} \quad (6)$$

where K_j designates all the respondents who selected interval j, q_k the certainty level of respondent k, and $n_{0.5}$, $n_{0.9}$ and $n_{0.99}$ are the number of respondents within K_j answering with the different certainty levels. For example, the physical environment weighted voting score for the interval $0,1\% - 0,5\%$ has been calculated simply as $w_j = 2 \cdot 0.5 + 1 \cdot 0.9 + 0 \cdot 0.99 = 1.9$. The weighted voting scores for physical environment are displayed graphically in Figure 2.

As a consequence of the distribution of the intervals, a linear plot is difficult to read. A logarithmic version is therefore given in Figure 2. To determine the probability p_i of a less than best practice physical environment factor to cause unavailability in enterprise IT systems, the interval with the highest weighted voting score is selected. To determine the exact location within the interval j, the weighted voting scores of the two adjacent intervals $j-1$ and $j+1$ are used, so that

$$p_i = j_{min} + \frac{w_{j+1}}{w_{j+1} + w_{j-1}} (j_{max} - j_{min}) \quad (7)$$

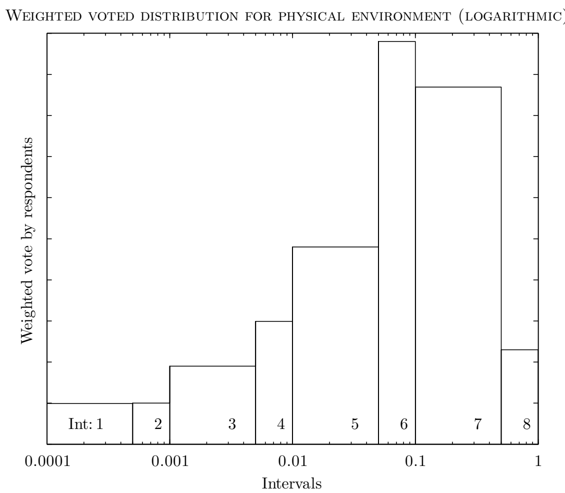

Figure 2. Weighted voting scores for physical environment, logarithmic version. $p_i \approx 0.082$ represents the share of currently unavailable enterprise IT systems that would, in the experts' opinion, be available if the physical environment had been managed according to best practice

where j_{min} and j_{max} designate the start and end points of interval j. In this case, as illustrated in Figure 2, the interval with the highest weighted voting score is number 6, $5\% - 10\%$, with $w_6 = 9.8$. The adjacent intervals have the weighted voting scores $w_5 = 4.8$ and $w_7 = 8.69$. These relative scores indicate that the probability p_1 should be located slighlty above the midpoint of the $5\% - 10\%$ interval. The calulation yields

$$p_1 = 0.05 + \frac{8.69}{8.69 + 4.8}(0.1 - 0.05) \approx 0.0822$$

The procedure is iterated for each and every causal factor, resulting in probabilities $p_1, \ldots, p_n$ as illustrated in Table III (rounded to one decimal). Each p_i reflects the share of currently unavailable enterprise IT systems that would, in the experts' opinion, be available if the *factor* X_i had been managed according to best practice. It might seem counterintuitive that $\sum p_i > 100\%$, but consider a system that went down because of an internal application error, and then did not come up because proper backups did not exist. At an appropriate time after the mishap, it is true that the system would have been available if the application error had been avoided, and also true that the system would have been available if the backups had been better. Thus, the factors need not be mutually exclusive.

As can be seen from Table III, judging from the respondents' answers, best practice change control is the factor most prone to increase availability of enterprise IT systems, closely followed by best practice component monitoring, and best practice requirements and procurement.

One factor is still missing in order to obtain a complete leaky Noisy-OR model, viz. the leakage p_0. To obtain an

Table III
CAUSAL FACTORS WITH PROBABILITIES FOR NOISY-OR MODEL.

	Causal factor X_i Lack of best practice …	p_i
1	…physical environment	8.2%
2	…requirements and procurement	25.2%
3	…operations	23.0%
4	…change control	28.1%
5	…technical solution of backup	7.0%
6	…process solution of backup	3.6%
7	…data redundancy	7.8%
8	…storage architecture redundancy	2.8%
9	…infrastructure redundancy	2.9%
10	…avoidance of internal application failures	16.9%
11	…avoidance of external services that fail	8.7%
12	…network redundancy	7.6%
13	…avoidance of network failures	18.3%
14	…physical location	3.3%
15	…resilient client/server solutions	3.6%
16	…monitoring of the relevant components	26.1%

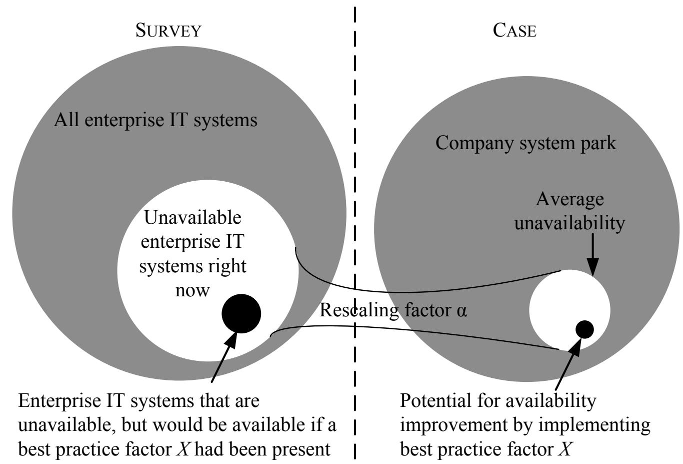

Figure 3. Venn diagrams schematically depicting the relation between the survey and an application case.

estimate of the leakage, the experts consulted in this survey was asked to comment if they believed that any important factors contributing to unavailability were left out in the survey. This is discussed in further detail in section VI. Suffice to note here that since no single proposed missing factor was mentioned by more than two experts (out of the 50 respondents) it seems safe to assume that the leakage should be less important than the least important factor considered in the survey. As seen in Table III, the smallest p_i belongs to storage architecture redundancy at 2.8%. A leakage $p_0 = 1\%$ therefore seems reasonable, and will be used throughout the remainder of the article.

C. Rescaling for case-based assessment

The questions answered by the respondents were explicitly concerned with increasing the availability of unavailable systems ("How large a share of currently unavailable enterprise IT systems would you guess would be available if a best practice *factor X* had been present?"). The leaky Noisy-OR model therefore explains enterprise IT systems *un*-availability (Y), employing the parameters $X_1, \ldots, X_n$ describing the lack of best practices, and the model is built in the domain of unavailable enterprise IT systems. However, a more practical typical concern is the availability of an entire park of systems, with a known prior availability baseline. The Bayesian model therefore needs to be rescaled from the set of unavailable enterprise IT systems to the whole set of enterprise IT systems. Figure 3 (slightly adapted from the survey) illustrates the issue.

Another way to express the issue is that the unscaled Noisy-OR model reflects the *potential for improvement*, by addressing only unavailability.

The most straightforward way to rescale the model, in order to answer how a system's availability can be improved by applying best practice solutions, is to apply a rescaling factor α to all p_i, including the leakage p_0. It could be argued that a single α should not be applied to all factors alike, but in the absence of good reasons to treat them separately, this is surely the simplest and best warranted solution. It follows from Equation (5) that

$$A(\mathbf{X}_p) = 1 - P(y|\mathbf{X}_p) = (1-\alpha p_0) \prod_{i:X_i \in \mathbf{X}_p} \frac{(1-\alpha p_i)}{(1-\alpha p_0)} \quad (8)$$

where $A(\mathbf{X}_p)$ is the availability of a given system lacking the best practice factors listed in the vector $\mathbf{X}_p$.

V. A SCENARIO-ANALYSIS EXAMPLE

This section is intended to illustrate how the leaky Noisy-OR model presented in the previous section can be used for actual assessment of an enterprise IT system, and how it can guide decision-making with an impact on availability.

To give an example, if one knows that the system *Saurischia* has a current availability of 99.8% and that best practice was only applied in the cases of data redundancy (X_7), storage architecture redundancy (X_8) and infrastructure redundancy (X_9), Equation (8) becomes

$$99.8\% = (1-\alpha p_0) \prod_{i=1}^{6} \frac{(1-\alpha p_i)}{(1-\alpha p_0)} \prod_{i=10}^{16} \frac{(1-\alpha p_i)}{(1-\alpha p_0)}$$

Solving for α (analytically this is cumbersome due to all the binomial coefficients, but numerically it is easy) yields a rescaling factor $\alpha \approx 0.00117223$.

Continuing the example of the *Saurischia* system, it is natural to ask how to improve its availability. Here, the model can give precious guidance. Assuming that the 13 factors not currently reaching best practice are at the decision-maker's disposal, their respective impacts can easily be analyzed and compared to the prior baseline of 99.8%.

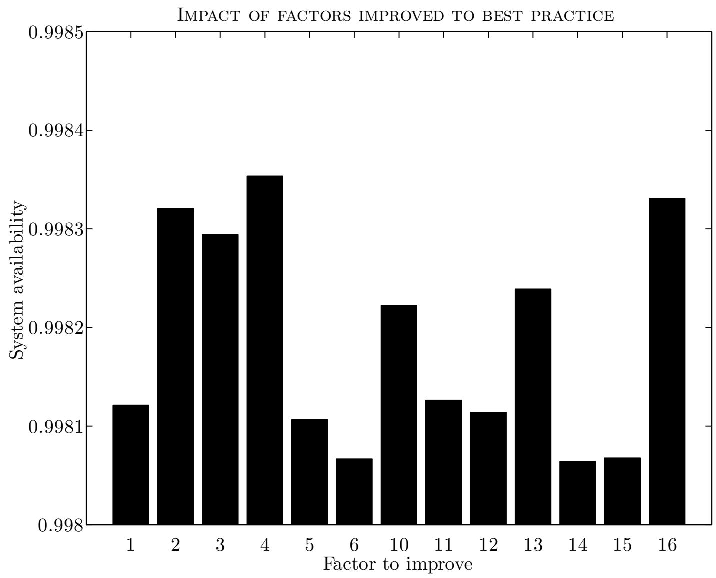

	Causal factor
1	Physical environment
2	Requirements and procurement
3	Operations
4	Change control
5	Technical solution of backup
6	Process solution of backup
10	Avoidance of internal application failures
11	Avoidance of external services that fail
12	Network redundancy
13	Avoidance of network failures
14	Physical location
15	Resilient client/server solutions
16	Monitoring of the relevant components

Figure 4. Prediction of how improvements of factors to the best practice-level would impact the availability of the example *Saurischia* system.

Figure 4 illustrates the predicted impact of each of these 13 factors taken by themselves. As can be seen, factor 4 (change control), factor 16 (monitoring of the relevant components), and factor 2 (requirements and procurement) are the most promising candidates for availability improvement of the *Saurischia* system. It might be objected that this could have been read straight off Table III – finding the most promising candidates requires only an ordinal ranking. However, a key strength of the Bayesian method is the possibility to investigate the impact of getting several factors up to best practice-level at the same time (as seen from Equation 8). To evaluate these interactions of several factors, Table III is not sufficient by itself, but the full leaky Noisy-OR model is needed. Another strength of the full model is that the expected cost of unavailability (e.g. from [1] or using a method like [8]) can be compared to the estimated costs for getting the various factors up to best practice-level.

VI. Discussion

A. The Noisy-OR assumptions

As noted in section III, proper use of the Noisy-OR gate makes two assumptions regarding the structure of the interaction of causes $(X_1, \ldots, X_n)$ and effect (Y) [18], [29]. These are (i) that each of the causes has a probability p_i of being sufficient for producing Y, and (ii) that the ability of each cause $X_1, \ldots, X_n$ to bring about Y is independent. In the present study, Y is unavailability of enterprise IT systems and $X_1, \ldots, X_n$ are causes of such unavailability.

Arguing for (i) in this case is straightforward. Indeed, it is almost always assumed that failing non-redundant components of complex systems can cause malfunctions by themselves. However, these faults are not always deterministic – e.g. a non-best practice requirements and procurement process will not infallibly lead to unavailability, but will do so with a certain probability p. Arguing for (ii) is harder. In many cases it is reasonable to assume that factors are independent, but this is not always the case. Backup systems, for example, only come into play when a system has failed and it is time to restore it. Therefore, the impact of factors such as *technical solution of backup* and the *process solution of backup* depend to some extent on other factors. In general terms, the distinction between proactive and reactive factors indicate that (ii) is an approximation that does not hold in all circumstances. In the full model, such dependencies could be modeled by rescaling different factors p_i with different factors α_i accounting for interactions. However, to accurately reflect these phenomena, more empirical data is needed. To conclude, while the assumptions required for the Noisy-OR model are reasonable as a first approximation, the model should certainly be open to further refinement. By and large, such refinement is a matter of empirical investigation, where availability data from enterprise IT systems can be analyzed and statistically checked for independence. It should be noted, however, that there is no need to refashion the entire Bayesian model should some cause variables $X_1, \ldots, X_n$ turn out to be dependent. The bulk of the causes can remain modeled in a Noisy-OR relation to each other, while a select few can be modeled using different CPDs.

B. Validity of the model

So far, the validity discussion has mainly focused on the numbers. As discussed in section III-B, the respondents were carefully selected based on scientific merit, the uncertainty of their answers was taken into account, and self-selection ensured that the 50 final respondent were among the most qualified. However, the discussion of leakage also leads to a discussion of the *completeness* of the model. First and foremost, a strong argument for the completeness of the model is that it is based on the widely cited [4]. However, converting the qualitative theory of [4] into questions suitable for building a quantitative Bayesian model unavoidably introduces distortions. Two questions thus need to be addressed: (i) are there causes missing that should be added? and (ii) are there superfluous causes that should be removed? Together, these questions determine whether the model contains all relevant causal factors.

(i) was explicitly addressed in the survey. The question "Do you believe that important aspects of enterprise IT systems availability have been left out in the survey? If so, please describe the areas missing." received 18 answers, i.e. 32 of the respondents did not find any aspects missing important enough to warrant an answer. Out of these 18 answers, two were in the negative, i.e. confirming the completeness of the model. Another two addressed methodology issues, that will be discussed below, but did not constitute suggestions for additional causal factors relevant to enterprise IT systems availability. The remaining 14 replies are summarized in Table IV.

Table IV
MISSING FACTORS IDENTIFIED BY 14 RESPONDENTS. SOME RESPONDENTS IDENTIFIED MORE THAN ONE MISSING FACTOR.

CAUSAL FACTOR	N
Security attacks	1
Automatic computing	1
Software quality	1
Disaster management	1
Human factors	2
Methodologies such as ISO-standards	1
Evolution/maintenance	1
Data availability and data model	1
Performance monitoring	1
Cost/benefit issues	2
Virtualization	1
Load management	1
Infrastructure should be more specific	1
Enterprise organization	1
Relationship between SLA and desired uptime	1

The two methodology questions raised were (a) that subjective perception of availability may differ from objective measures, and (b) that practitioners rather than academics should have been selected as respondents. These are both relevant points, but they are also complementary in an interesting way. Asking practitioners to give estimates would run a higher risk of being influenced by subjective perceptions (since a user or administrator dealing with systems on an everyday basis has the opportunity to develop a subjective perception, as opposed to a scientist collecting data or building models in a fashion more disconnected from daily system usage). Conversely, asking published scientists limits the risk of subjective perceptions based only on ones own systems (since scientific publication requires a certain generality, and a careful discussion of validity), but at the same time runs the risk of missing valuable "down-to-earth" insights from the practitioner community. There seems, thus, to exist an inherent methodological trade-off between (a) and (b), and in the light of this, receiving one comment on each is not a bad result. The details on the selection of survey participants was more thoroughly discussed in section III-B.

As seen in Table IV, no single potentially missing causal factor was identified by more than 4% of survey respondents (2 out of 50). Most were identified by just a single respondent. Since there is no strong agreement among the 50 experts on which causal factors are missing, we conclude that the model contains an appropriate set of causal factors causing enterprise IT system unavailability.

As for concern (ii) – superfluous causes of unavailability in the model that should be removed – it was not explicitly addressed in the survey in the sense that any particular question was devoted to it. However, every question on causes implicitly addresses the issue, as the respondents could always say that a very minute fraction ($< 0,05\%$) of currently unavailable enterprise IT systems would be available if a best practice *factor X* had been present. It should be noted, of course, that this is not an unambiguous measure of the superfluousness of a cause. A causal factor that is both very important to availability and very well managed in the real world does not offer the kind of potential for improvement that the question looks for. However, as discussed in section IV, it does offer a measure of the practical relevance of the causal factor. A causal factor with a large potential for improvement is, ceteris paribus, more relevant to a practitioner than a causal factor with a small potential for improvement.

VII. CONCLUSIONS

The contribution of the present paper is three-fold. First, the results from an academic expert survey on the causes of unavailability of enterprise IT systems are presented. Second, these results are used to build a Bayesian decision support model for assessment of enterprise IT systems availability. Third, an example is presented to illustrate how the model can actually be put to use by practitioners aiming to ensure systems availability.

A natural continuation of the present line of research is to validate the results with case studies of actual enterprise IT systems. Empirical data from such investigations could be used both quantitatively – to calibrate the numbers in the Bayesian model, and qualitatively – to restructure the Bayesian network if the leaky Noisy-OR assumptions prove unsuitable for some variables. Another line of prospective future work is to embed the present model into a larger enterprise architecture framework for IT systems availability analysis, along the lines of [12].

Acknowledgment

The authors are grateful to those colleagues and respondents who helped improving the survey before it was sent out, as well as to all the researchers who responded when invited. A special thanks is due to Lars Nordström for his contribution and support in the early part of the project.

References

[1] IBM Global Services, "Improving systems availability," IBM Global Services, Tech. Rep., 1998.

[2] J. Askåker and M. Kulle, "Miljardaffärer gick förlorade," *Dagens Industri*, pp. 6–7, June 4 2008.

[3] B. Malik, "Q&A: How Much Does an Hour of Downtime Cost?" Gartner, Inc., Tech. Rep., Sep. 2009.

[4] E. Marcus and H. Stern, *Blueprints for high availability, second edition*. Indianapolis, IN, USA: John Wiley & Sons, Inc., 2003.

[5] K. Campbell, L. A. Gordon, M. P. Loeb, and L. Zhou, "The economic cost of publicly announced information security breaches: empirical evidence from the stock market," *J. Comput. Secur.*, vol. 11, no. 3, pp. 431–448, 2003.

[6] International Organization for Standardization, "Software engineering - product quality - part 2: External metrics," International Organization for Standardization, International standard ISO/IEC TR 9126-2:2003(E), July 2003.

[7] J. Saat, U. Franke, R. Lagerström, M. Ekstedt, and R. Winter, "Meta modeling for situational IT/business alignment," 2010, submitted manuscript.

[8] D. Scott, "How to Assess Your IT Service Availability Levels," Gartner, Inc., Tech. Rep., Apr. 2009.

[9] M. Malek, B. Milic, and N. Milanovic, "Analytical availability assessment of it services," in *ISAS*, 2008, pp. 207–224.

[10] N. Milanovic, B. Milic, and M. Malek, "Modeling business process availability," in *SERVICES '08: Proceedings of the 2008 IEEE Congress on Services - Part I*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 315–321.

[11] X. Zhang and H. Pham, "An analysis of factors affecting software reliability," *J. Syst. Softw.*, vol. 50, no. 1, pp. 43–56, 2000.

[12] R. Lagerström, U. Franke, P. Johnson, and J. Ullberg, "A method for creating enterprise architecture metamodels - applied to systems modifiability analysis," *International Journal of Computer Science & Applications*, vol. VI, pp. 89–120, Dec. 2009.

[13] R. Lagerström, P. Johnson, D. Höök, and J. König, "Software Change Project Cost Estimation – A Bayesian Network and a Method for Expert Elicitation," in *Third International Workshop on Software Quality and Maintainability (SQM 2009)*, Mar. 2009.

[14] N. Friedman, M. Linial, and I. Nachman, "Using bayesian networks to analyze expression data," *Journal of Computational Biology*, vol. 7, pp. 601–620, 2000.

[15] R. E. Neapolitan, *Learning Bayesian Networks*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 2003.

[16] F. V. Jensen, *Bayesian Networks and Decision Graphs*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2001.

[17] R. D. Shachter, "Probabilistic inference and influence diagrams," *Oper. Res.*, vol. 36, no. 4, pp. 589–604, 1988.

[18] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.

[19] P. H. Garthwaite, J. B. Kadane, and A. O'Hagan, "Statistical methods for eliciting probability distributions," *Journal of the American Statistical Association*, vol. 100, pp. 680–701, June 2005. [Online]. Available: http://ideas.repec.org/a/bes/jnlasa/v100y2005p680-701.html

[20] F. Johansson and G. Falkman, "Implementation and integration of a bayesian network for prediction of tactical intention into a ground target simulator," in *Information Fusion, 2006 9th International Conference on*, July 2006, pp. 1–7.

[21] S. Renooij, "Qualitative approaches to quantifying probabilistic networks," Ph.D. dissertation, Utrecht University, Utrecht, the Netherlands, 2001.

[22] T. W. Mangione, *Mail surveys. Improving the quality*. Sage Publications Inc., 1995.

[23] F. J. J. Fowler, *Survey research methods*. Sage Publications Inc., 2002.

[24] L. Blaxter, C. Hughes, and M. Tight, *How to research*. Open University Press, 2006.

[25] R. Czaja and J. Blair, *Designing surveys*. Sage Publications Inc., 2005.

[26] G. Baecher, "Judgemental probability in geotechnical risk assessment," The Office of the Chief, U.S. Army Corps of Engineers, Tech. Rep., 1988.

[27] O. Woodberry, A. E. Nicholson, K. B. Korb, and C. Pollino, "Parameterising bayesian networks," in *AI 2004: Advances in Artificial Intelligence*. Springer Berlin / Heidelberg, 2005, pp. 1101–1107.

[28] R. M. Cooke, *Experts in Uncertainty. Opinion and Subjective Probability in Science*. Oxford University Press, 1991.

[29] "Learning bayesian network parameters from small data sets: application of noisy-or gates," *International Journal of Approximate Reasoning*, vol. 27, no. 2, pp. 165 – 182, 2001.

[30] N. Friedman and M. Goldszmidt, "Learning Bayesian networks with local structure," in *Graphical Models*, M. I. Jordan, Ed. Cambridge, MA: MIT Press, 1999, pp. 421–459.

[31] M. Henrion, "Some practical issues in constructing belief networks," in *Uncertainty in Artificial Intelligence 3*, L. Kanal, T. Levitt, and J. Lemmer, Eds. North Holland: Elsevier Science Publishers B.V., 1989, pp. 161–173.