Balancing safety and availability for an electronic protection system

S. Wagner, I. Eusgeld & W. Kröger
Laboratory for Safety Analysis, ETH Zurich, Zurich, Switzerland

G. Guaglio
MEMC Electronic Materials, Novara, Italy

ABSTRACT: A generic methodology is being developed for addressing the trade-off between safety and availability related to large and complex electronic protection systems. The methodology allows the modelling of a highly fault tolerant and safe system, using an object-oriented approach to build a model frame which merges Monte Carlo method and results of Fault Tree Analysis. This paper introduces the methodology and demonstrates its feasibility and suitability by means of a case study performed on the Machine Protection System (MPS) of the Large Hadron Collider (LHC) at CERN, the European Organization for Nuclear Research.

ABBREVIATIONS

BEE	Back-end electronics board
BIC	Beam interlock controller
BIS	Beam Interlock System
BLMS	Beam Loss Monitor System
FEE	Front-end electronics board
FTA	Fault Tree Analysis
IC	Ionisation chamber
LBDS	Beam Dumping System
LHC	Large Hadron Collider
MPS	Machine Protection System

1 INTRODUCTION

The trade-off between safety and availability of technical equipment is one of the main issues for protection systems. While ensuring safe operation by triggering shutdowns in case of dangerous equipment conditions, a protection system should not cause operation interruptions due to false alarms. In case of the LHC MPS, 'safety' means LHC operation under defined conditions and the protection of an investment of about 3 billion Euro, respectively. 'Availability' refers to beam availability, i.e. the LHC providing particle beams for the experiments.

Several studies have been made on the LHC MPS addressing this trade-off (Filippini 2006; Guaglio 2005; Todd 2006; Vergara Fernandez 2003). They mainly focus on individual MPS subsystems and base upon classical methods such as Failure Mode and Effects and Criticality Analysis (Todd 2006), FTA (Guaglio 2005) and Markov models (Filippini 2006; Vergara Fernandez 2003). Based on these previous studies, the present approach aims at a global analysis covering the whole MPS. Due to the size and complexity of the MPS, this aim is regarded as hardly achievable with classic techniques exclusively. A new modelling concept for a highly fault-tolerant and safe system with respect to the trade-off between availability and safety is required. For this purpose, a generic methodology is being developed. It uses an object-oriented approach to build a model frame which merges Monte Carlo method with results of FTA.

While Monte Carlo simulation is a common approach, object-oriented modelling is not widely-used for system reliability analysis yet. However, the potential of an object-oriented modelling approach in this field has been shown, e. g. for the optimization of maintenance strategy (Kaegi and Mock 2007) and the reliability (availability) analysis and optimization of a traffic network (Bonabeau 2002).

The combination of event tree analysis and Monte Carlo-based simulation with concepts from object-oriented analysis has been successfully applied to an aviation safety problem (Wyss et al. 2004). Another example of an object-oriented approach involving Monte Carlo simulation is the reliability analysis of electric power systems applied by Schläpfer et al. (2008).

The developed methodology differs from these approaches by explicitly including FTA. Furthermore, it distinguishes itself through the application to a large and complex electronic protection system. The methodology provides a straightforward bottom-up modelling and simulation approach for this type of systems. This paper includes an evaluation of the methodology based on first experience with its application to the LHC MPS. The focus lies on two main aspects with regard to availability: 1) the proportion of shutdowns due to false alarms and 2) the importance of the related components.

2 METHODOLOGY

This section describes the developed methodology in a generic way. Its application to the LHC MPS is presented in the next section. The term 'system' hereafter refers to 'electronic protection system'.

2.1 *Requirements*

The trade-off between safety and availability of a system is reflected in the borderline between fault tolerant and fail-safe design. A fault tolerant design increases the reliability of the system and thereby contributes to the safety of the protected equipment. In case of non-tolerated failures within the system, fail-safe measures assure that the protected equipment is turned into a safe state, which most often corresponds to a shutdown of operation. However, a too restrictive application of the fail-safe principle causes unnecessary shutdowns. Therefore, a methodology addressing that trade-off must involve redundancy and self-monitoring, which are main elements in fault tolerant and fail-safe design, respectively. Additionally, the system demand, i.e. the occurrence of dangerous conditions of the protected equipment requiring the system's action, has to be taken into account. The multiple failure modes of electronic devices also need to be dealt with.

The methodology is expected to give insight into the system's global behaviour, focusing on the following parameters:

- Probability of missed action upon demand (affecting safety)
- Probability of false alarms (affecting availability)
- Critical components in terms of missed action and false alarms
- Contribution of redundancy and fail-safe measures to the system's behaviour

2.2 *Hybrid concept*

The underlying concept of the developed methodology is illustrated in Figure 1. An object-oriented model approach builds the frame of the methodology. The model basically reproduces the primary signal path of the system, i.e. following (IEC 1998), the signal path from the sensors via a sequence of other components to the actuators. Individual FTA of the components in the path provide component's reliability numbers (failure rates), which form the basis for the Monte Carlo simulation of the system's global behaviour.

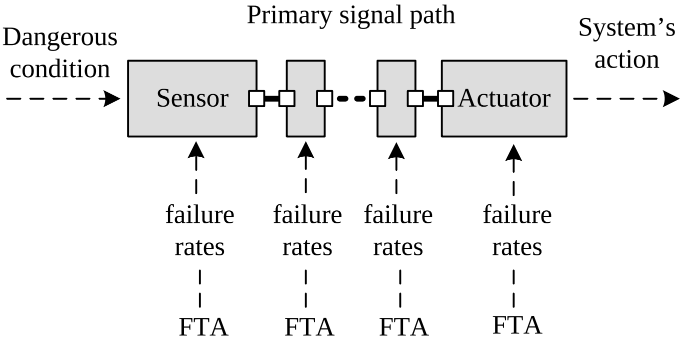

Figure 1: Concept of the methodology

The major advantage of such an object-oriented model is given by it's bottom-up approach. The system can basically be modelled component by component, following the primary signal path. In-depth knowledge of the system's global behaviour is not necessary for modelling, but emerges from the simulations.

2.2.1 *Global object-oriented model of the system*

The components of the primary signal path are modeled as individual objects that are linked according to the physical system. The objects are designed in a simple 'black box' manner (Fig. 2). The component behaviour is described by means of three states (represented by a Markov model). Two different failure modes are distinguished, referred to as 'blind' and 'false', while the failure-free state is called 'ready':

- **ready** An incoming alarm signal passes the object and is transmitted to the following object. This state corresponds to the component's full functional capability.
- **blind** An incoming signal is not transmitted, the signal path is cut. This state refers to failures ('blind failures') that inhibit the component's function.
- **false** A new signal is generated and transmitted, independently from incoming signals. This state relates to failures ('false failures') that are covered by fail-safe measures and lead to a false

alarm, i.e. self-triggering of the system in absence of dangerous equipment conditions.

The failures of the components are assumed to be independent. Simultaneous occurrence of such failures is included in the model. This is of special interest if dangerous combinations arise, e.g. if an alarm signal due to a failure in the equipment (demanding the system's action for shutdown) meets a 'blind' component.

This generic object model is applied to all the components. The individual models of the components only differ in terms of number of input and output gates and state probabilities and transition rates, respectively. The transition rates are given by the results of the component's individual FTA (Section 2.2.2).

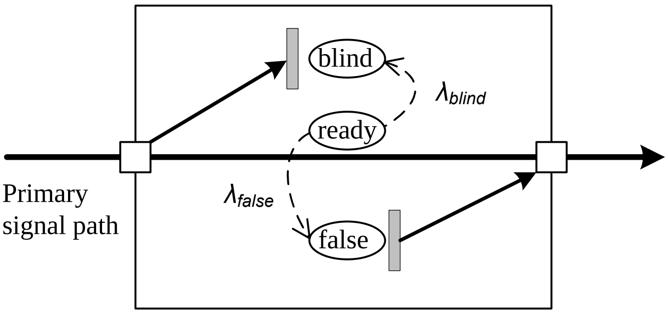

Figure 2: Basic component model

The described model includes the following assumptions:

- Fail-safe measures leading to false alarms are inherent to the components and independent, i.e. there's no common additional self-monitoring system involved contributing to common cause failures.
- 'False' is assigned to the component that triggers the false alarm, which does not necessarily correspond to the location of the failure occurrence

Since signals are modelled as alarm signals, the model implies the forestalling of signal threshold comparison that may only take place further down the path.

2.2.2 *FTA of the components*

The data for the state transitions of the components result from FTA with top events defined as 'blind' and 'false'. The interfacing data between FTA and the object-oriented model frame are either represented by failure rates or by Weibull parameters. The first applies if the failure probability resulting at the top level of the fault tree follows an exponential distribution function. This is the case if constant failure rates underlie the fault tree and only disjunct events are involved. Weibull parameters are applied if the resulting failure probability is to be approximated by a cumulative Weibull distribution function. This is the case in a fault tree with conjunct events, e.g. due to redundancy within the component.

FTA is an established approach for the reliability analysis of electronic components (Schneeweiss 1999). In the introduced hybrid concept, the FTA is limited to the component level which keeps it easily manageable and traceable. The fault tree can be developed to any level of detail, ensuring accurate interface data at the top level. Since the interface data define independent state transitions in the model, independence of the fault trees must be assured.

2.2.3 *Stochastic simulation*

Following the questions of interest of the methodology (Section 2.1), the time to a missed action upon demand and a shutdown due to a false alarm respectively is simulated, taking into account operational cycles of the system. The demand of the system is included by either stochastic or deterministic generation of alarm signals at the beginning of the signal path (Fig. 1).

The simulation log file contains the following data:

- Occurred event
- Time of occurrence
- Related component

with the different types of events being

- **System demand** upon signal generation at the beginning of the signal path
- **False alarm generated** upon transition from 'ready' to 'false'
- **Component turned blind** upon transition from 'ready' to 'blind'
- **Signal not transmitted** upon signal meeting blind component
- **System action fulfilled** upon arrival of signal in actuator at the end of the path

3 APPLICATION TO THE LHC MPS

The Large Hadron Collider (LHC) at CERN is a particle accelerator that will collide two counter-rotating proton beams (beam 1 and beam 2) of very high intensities (Schmidt et al. 2006). The energy stored in the beams and the magnet system requires a complex machine protection represented by the MPS.

3.1 *The LHC MPS*

The primary MPS task is the protection of the LHC equipment against damage due to uncontrolled beam loss (Schmidt et al. 2006). This task includes the following basic functions:

1. Detection of dangerous conditions of beams or LHC equipment and generation of beam dump request signals
2. Transmission and concentration of the dump requests
3. Subsequent extraction of the beams from the accelerator into absorbers (beam dump blocks)

Corresponding to these functions, the MPS is made up by the following main subsystems:

- **Monitoring Systems** monitor the beams or the equipment and generate dump requests in case of dangerous conditions
- **Beam Interlock System (BIS)** transmits the dump request from the monitoring systems to the Beam Dumping System
- **Beam Dumping System (LBDS)** realizes the extraction of the beams from the accelerator and their disposal in the absorbers

The primary signal path (Section 2.2.1) starts in the monitors (sensors) of the Monitoring systems, continues through the BIS and ends at the magnets (actuators) of the LBDS.

3.2 *The model*

The current model includes the Beam Loss Monitor System (BLMS) as a monitoring system and the BIS. The LBDS is assumed to be fault-free.

Modelling and simulation have been implemented using the commercial software AnyLogic 5 (XJ Technologies 2005).

3.2.1 *Global model structure*

The level of detail of the MPS on which the components are treated as 'black box' (as presented in Section 2.2.1) roughly corresponds to the level of electronic circuit boards. The resulting global structure of the model is illustrated in Figures 3 and 4. Figure 3 shows one of eight branches of the BLMS, with terminology following Guaglio (2005). The signal path starts in the ionisation chambers (IC). The signals of six ICs are treated by one front-end electronics board (FEE). The back-end electronics board (BEE) treats the signals of two FEEs. Thirteen daisy-chained BEEs and one Combiner Card build a VME crate. Three VME crates are daisy-chained and transmit the signal to one CIBU-S, which is the interface component to the BIS.

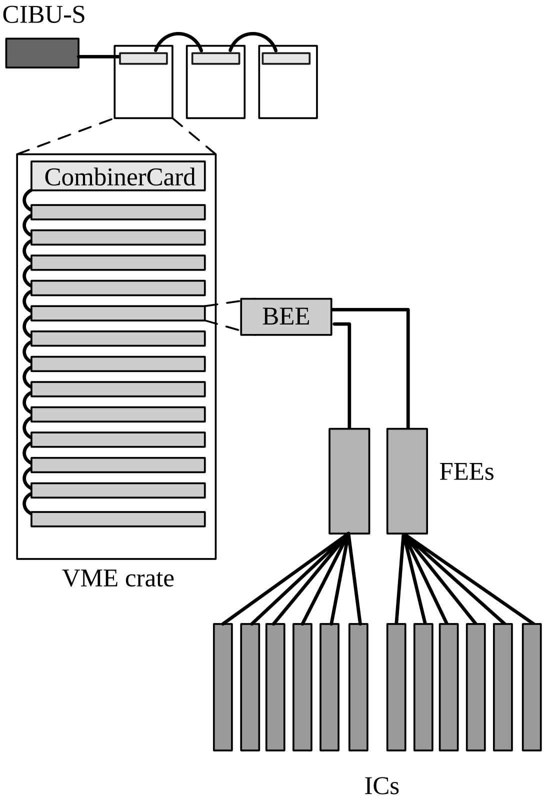

Figure 3: Structure of one BLMS branch

.

The BIS (Fig. 4) consists of thirty-two beam interlock controllers (BIC). Sixteen of them treat signals which relate to beam 1, the other half treats signals relating to beam 2. The BICs of each beam are daisy-chained with two counter-rotating redundant signals paths, i.e. the signals coming from the CIBU-S are transmitted to both neighboring BICs, ending at the interface components to the LBDS. Since the signals from the BLMS apply to both beams, one CIBU-S is linked to both BICbeam1 and BICbeam2.

The resulting numbers of components included in the model is given in Table 1. Both the structure of the model and the number of components reflect the MPS.

3.2.2 *Component model*

The implementation of the component model in AnyLogic directly corresponds to the approach presented in Section 2.2.1. The component is modelled as an object which is graphically represented by a box with the number of input and output gates corresponding to the primary signal path (Fig. 5, top). The component behaviour is included by an encapsulated object called

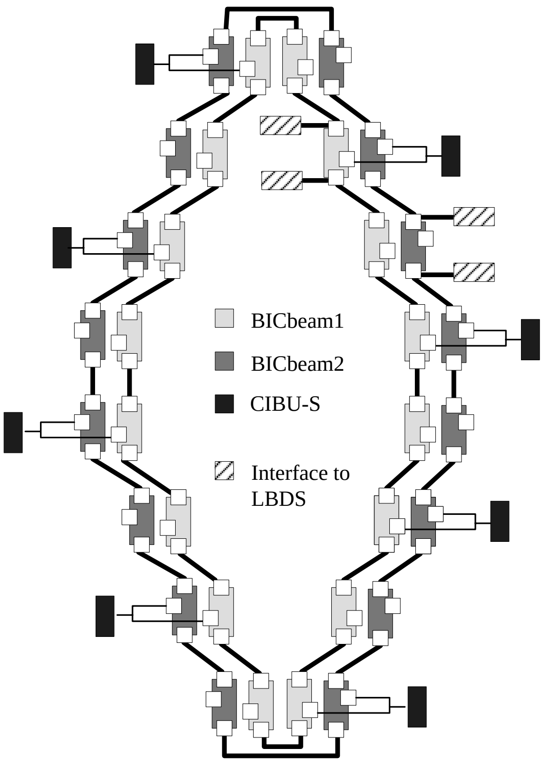

Figure 4: Structure of BIS

Table 1: Number of components

i	Component i	Number n_i
1	IC	3744
2	FEE	624
3	BEE	312
4	Combiner Card	24
5	VME crate	24
6	CIBU-S	8
7	BICbeam1	16
8	BICbeam2	16
	Total	4768

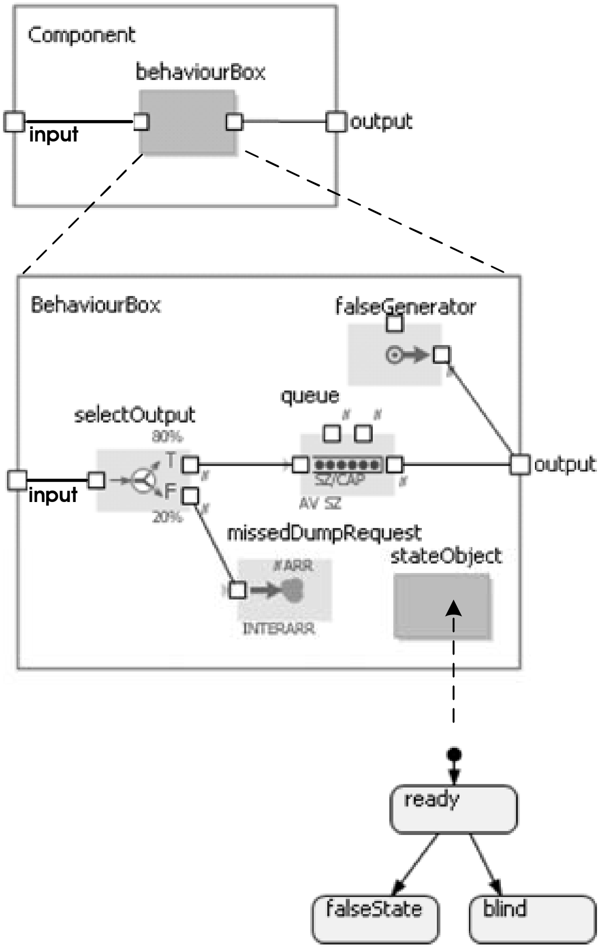

Figure 5: 'Black box' component model implemented using AnyLogic

BehaviourBox corresponding to the state model introduced in Figure 2. The elements used in the signal path of the BehaviourBox are objects taken from AnyLogic Enterprise Library.

The signal path is defined by the state diagram placed in the object StateObject (Fig. 5, bottom). Upon transition to 'falseState', a new signal entity is generated in the FalseGenerator object and transmitted to the output gate. Upon transition to 'blind', the condition for the SelectOutput object switches from true (T) to false (F). Incoming signals are lead to the MissedDumpRequest object, which is a dead end for the signal.

The use of this generic component model implies the following assumptions:

- Redundancy of lines between two components is not reproduced in the model. The behaviour of the lines is assigned to the subsequent component and is taken into account in the FTA of that component.
- Signal transmission is assumed to be immediate, i.e. transmission delay is not included. This seems justifiable in view of MPS transmission delays in the range of milliseconds and LHC operational cycles in the range of hours (Schmidt et al. 2006).
- The failures of the components are independent, common cause failures are not taken into account.
- Maintenance is not taken into account.

3.3 *Input data*

The input data of the model includes component's state transition data based on FTA and signal generation data for the ICs, relating to the demand of the system.

3.3.1 *Component's failure rates*

The failure rates used for the present simulations are given in Table 2. They represent orders of magnitudes based on the study performed by Guaglio (2005). His work includes fault trees of a BLMS branch with top events corresponding to 'blind failure' and 'false failure' of the BLMS.

Table 2: Failure rates of the components

i	Component i	Rate λ_i^{false}	Rate λ_i^{blind}
		1/h	1/h
1	IC	1E-7	1E-7
2	FEE	1E-6	1E-8
3	BEE	1E-8	1E-9
4	Combiner Card	1E-8	1E-9
5	VME crate	1E-5	1E-8
6	CIBU-S	1E-6	1E-13
7	BICbeam1	1E-5	1E-13
8	BICbeam2	1E-5	1E-13

As a first approximation and following the rare event approach described by Guaglio (2005), the (constant) failure rates of the basic events in the fault trees were summed up to the level corresponding to the components in the model. The resulting order of magnitude are taken as constant failure rates defining exponentially distributed times to state transition in the model. The rates for the BIS components base upon expert judgement (Todd, pers. comm.).

MPS self-monitoring involves a series of testing procedures, here exemplified by the BLMS. In the BLMS, four testing intervals are present: two seconds, two minutes, LHC cycle and year (Guaglio 2005). The former two relate to continuous tests within the mission, the latter two to tests taking place in between missions (i.e. with no beams in the LHC). The tests with two second and two minute intervals contribute to the 'false failures'. If a test detects a failure in the BLMS, a false dump is triggered (by definition, false dumps refer to dumps triggered by the MPS in absence of dangerous LHC equipment conditions, Section 2.2.1). The 'blind failures' include failure modes not covered by tests. This is the case if 1) there is no test implemented for the failure mode, 2) the system is only tested before each new LHC cycle(i.e. in between the missions) or annually or 3) the continuous testing fails with the failure left 'blind'.

Given the immediate transmission assumption (Section 3.2.2), the eventual (blind) latency of two minutes in the execution of the false alarm upon failure detection is neglected in the model. This does not corrupt the 'blind' calculation, because the latency is covered by case 3) mentioned above.

3.3.2 *System demand*

In the first simulations described in this paper, the system demand is modeled by a timer that expires following an exponential distribution function with a Mean Time to Failure ($MTTF^{beamLoss}$) of 100 hours. Upon expiry, one of the seventy-eight FEEs of the first BLMS branch is chosen following a discrete uniform distribution function, and a signal is generated in the six related ICs.

The described model of system demand represents a simplified approach for these first simulations. In reality, beam loss occurs anywhere in the LHC, not only in one sector, i.e. one branch of the BLMS. Furthermore, a beam loss event is not locally limited to the six ICs covering the area of a magnet but spreads over more extended areas of the machine, triggering several ICs at different locations. A demand model corresponding to the LHC beam loss pattern is to be implemented in a next step.

3.4 *Simulation specifications*

A nominal LHC operational cycle is expected to last about twelve hours and basically includes injection of the beams, ramping up the beams and 'physics', which names the cycle phase at top beam energy where the experiments at the collision points run (Schmidt et al. 2006). At the end of physics, the mission is ended by a scheduled end-of-mission dump and the LHC is prepared for a new cycle. An early end of a mission is caused either by emergency dumps or by false dumps. The first relate to dumps triggered due to dangerous beam or LHC equipment conditions. The latter refer to dumps triggered by failures within the MPS.

In accordance with the LHC operational cycle, a simulation run is stopped after twelve hours of model time or upon arrival of a signal entity in the interface components to the LBDS at the end of the signal path. In the simulation stop after twelve hours of model time, a fault-free reaction of the system to an end-of-mission dump request is implied. Furthermore, twelve hours of mission time are slightly too conservative an assumption, since they also embrace the ramping down and pre-injection phase after a dump, where there is no beam in the LHC. The start of the simulation run corresponds to the start of beam injection into the LHC. Each start of a simulation run implies the model as-good-as-new.

According to Section 2.2.3, the simulation log includes event, time and component. The events, specified to the MPS, are the following:

- **beam loss** corresponding to system demand
- **false dump request** corresponding to false alarm generation
- **blind** corresponding to component turning blind
- **missed dump request** corresponding to signal not transmitted by component

- **dump triggered** corresponding to system action fulfilled

The worst case, the complete missing of an emergency dump results from 'missed dump requests' by all (redundant) components of the related signal path. In the simulation, this case is implied in a missing 'dump triggered' event upon occurrence of a 'beam loss' event.

3.5 *Results*

The results presented here base upon a first series of simulations including 100,000 missions. The simulations were performed on an Intel SC5299 Server with two Quad-Core 5335 Processors (2 GHz) and 8 GB of RAM, running Linux. The simulation time came up to approximately 12 hours.

Figure 6 gives the fraction of early ended missions due to emergency dumps and false dumps. In total, 11,277 or 11.3% of the missions ended with an emergency dump triggered by a beam loss event, 1749 or 1.7% of the missions with a false dump due to a false dump request of a component.

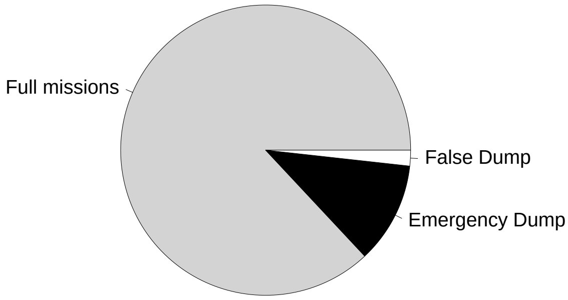

Figure 6: Fraction of early ended missions due to emergency or false dumps

The total fraction of early ended missions emanated from the simulations is verified by the cumulative exponential distribution function

$$F(t) = 1 - e^{-\lambda^{tot} \cdot t}$$

where

$$\lambda^{tot} = \sum_{i=1}^{8} \left(n_i \cdot \lambda_i^{false} \right) + \lambda^{beamLoss}$$

and t = 12 h; n_i = number of component i (Table 1); λ_i^{false} = 'false failure' rate of component i (Table 2); $\lambda^{beamLoss}$ = rate of beam loss ($1/MTTF^{beamLoss}$). The relative error comes up to 0.5%. The occurrence of the (comparatively) rare blind events is verified accordingly, resulting in a higher relative error of 2%.

With regard to balancing machine safety and beam availability, the identification of critical components and their contribution to the system's behaviour is essential. Figure 7 shows the contribution of the components to false dumps by triggering false dump requests. Almost 40% of the false dumps are caused by FEEs, while CIBU-S, Combiner Card and BEE scarcely contribute.

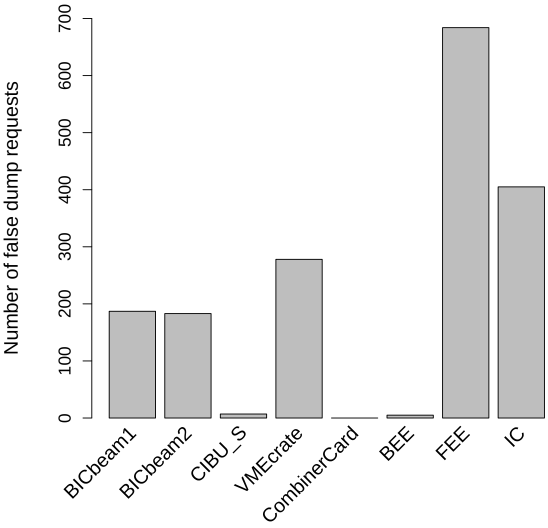

Figure 7: Components' contribution to false dumps

The analysis for the components' contribution to blind events unfolds an importance of almost 98% for ICs (447 blind events in total), which seems reasonable in view of the amount of ICs and their failure rates.

The combination of blind component meeting a demand results in a missed dump request. The present amount of simulation data is too limited to allow a contribution analysis with regard to missed dump requests (no missed dump request event in 100,000 missions). A rare event approach is to be implemented to the model in order to address these safety influencing events.

An overlap of many blind components together with a demand leads to the worst case scenario of a missed emergency dump. In view of extensive redundancy in the system, this scenario is expected to occur mainly upon common cause failures affecting large parts of the system. The inclusion of common cause failures is one of the next steps of model extension.

4 CONCLUSIONS AND OUTLOOK

The paper presents a new modelling concept for a highly fault tolerant and safe system with respect to the trade-off between availability and safety. The developed methodology uses an object-oriented approach to build a model frame which merges Monte Carlo method and results of FTA. Due to a generic 'black box' model of the system's components, the approach is straightforward and easily traceable, which is a crucial issue in the analysis of large and complex systems. Another major advantage is the underlying bottom-up principle, which distinguishes the approach from classic methods. For the setup of the model, in-depth knowledge of the global system's behaviour is not required. Detailed knowledge is only needed on component's level where it is manageable using fault trees. The approach allows addressing both safety and availability aspects by means of the same model, since all different failure modes of the components are included. The inclusion of the system demand, i.e. beam loss, involves the weighting of different system branches as an additional feature.

The methodology's application to the LHC MPS has validated its feasibility and suitability. For first results related to the LHC availability, the proportion of early ended missions due to false dumps (caused by failures of MPS components) and the importance of the related components have been shown and discussed.

The development of the methodology is going on with the expansion of the model to the LBDS and the inclusion of demand upon end-of-mission dump request. In a next step, a rare event approach is to be implemented, as well as a feature for including common cause failures. The testing and maintenance between the missions is an additional entry for further model extension. The upcoming operational experience with LHC and MPS provides the basis for the validation of the methodology and needs to be factored in the further development of the model.

ACKNOWLEDGEMENT

The authors thank all colleagues at CERN who have been supporting the work presented in this paper, in particular Rüdiger Schmidt, Jörg Wenninger, Bernd Dehning, Benjamin Todd, Markus Zerlauth and Jan Uythoven.

REFERENCES

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America 99*, 7280–7287.

Filippini, R. (2006). *Dependability analysis of a safety critical system: The LHC Beam Dumping System at CERN*. Ph. D. thesis, Università di Pisa.

Guaglio, G. (2005). *Reliability of the Beam Loss Monitors System for the Large Hadron Collider at CERN*. Ph. D. thesis, Université Blaise Pascal.

IEC (1998). Functional safety of electrical/ electronic/ programmable electronic safety-related systems (IEC 61508-4). IEC.

Kaegi, M. and R. Mock (2007). Agent-based simulation of maintenance for system optimisation. In *Proceedings of the European Safety and Reliability Conference 2007 (ESREL 2007)*, Stavanger, Norway.

Schläpfer, M., T. Kessler, and W. Kröger (2008). Reliability analysis of electric power systems using an object-oriented hybrid modeling approach. In *Power Systems Computation Conference 2008 (PSCC 2008), accepted*, Glasgow, Scotland.

Schmidt, R., R. Assmann, E. Carlier, B. Dehning, R. Denz, B. Goddard, E. B. Holzer, V. Kain, B. Puccio, B. Todd, J. Uythoven, J. Wenninger, and M. Zerlauth (2006). Protection of the cern large hadron collider. *New Journal of Physics 8*, 290. 31 p.

Schneeweiss, W. G. (1999). *The Fault Tree Method*. LiLoLe.

Todd, B. (2006). *A beam interlock system for CERN high energy accelerators*. Ph. D. thesis, Brunel University West London.

Vergara Fernandez, A. (2003). *Reliability of the Quench Protection System for the LHC superconducting elements*. Ph. D. thesis, Universitat Politècnica de Catalunya.

Wyss, G. D., F. A. Duran, and V. J. Dandini (2004). An object-oriented approach to risk and reliability analysis: Methodology and aviation dafety applications. *Simulation 80*, 33–43.

XJ Technologies (2005). *Anylogic user's manual*. XJ Technologies.