

Reliability Considerations and Fault Handling Strategies for Multi-MW Modular Drive Systems

Tobias Geyer, *Member, IEEE* and Stefan Schröder, *Member, IEEE*

Abstract—Shunt-interleaved electrical drive systems consisting of several parallel medium-voltage back to back converters enable power ratings of tens of MVA, low current distortions and a very smooth airgap torque. In order to meet stringent reliability and availability goals despite the large parts count, the modularity of the drive system needs to be exploited and a suitable fault handling strategy that allows the exclusion and isolation of faulted threads is required. This avoids the shutdown of the complete system and enables the drive system to continue operating. If full power capability is also required in degraded mode operation, redundancy on a thread level needs to be added. Experimental results confirm that thread exclusion allows the isolation of the majority of faults without affecting the mechanical load. As the drive system continues to run, faulted threads can be repaired and then added on-the-fly to the running system by thread inclusion. As a result, the downtime of such a modular drive system is expected to not exceed a few hours per year.

Index Terms—Reliability; fault handling; redundancy; medium-voltage drive; electrical drive; modular drive system; multilevel converter

I. INTRODUCTION

One way to achieve very high power levels in the range of tens of MVA is to operate several back to back converters (threads) in parallel. A de facto industrial standard for such a thread is the three-level medium-voltage Neutral Point Clamped (NPC) converter using Integrated Gate Commutated Thyristors (IGCT) as switching devices. Such a thread is usually rated at 8 to 12 MVA. Often, the N parallel threads are tightly coupled with very small inductances and the same switching pattern is applied to each of the threads. Doing this, the machine voltage resembles a three-level NPC inverter with effectively N times the thread power rating [5].

In many applications, however, a very smooth machine current and airgap torque is required. This can be achieved by coupling the threads with larger inductors (i.e. coupling inductors in the range of a quarter pu referred to the thread). This allows one to apply thread specific switching patterns that are optimized such that the resulting average switching patterns at the points of common coupling (the machine or the grid) exhibit significantly finer steps and resemble a higher switching frequency. One way to achieve this is to shift the PWM carrier signal by a certain phase shift among the threads giving raise to the concept of interleaving, see e.g. [6] and [17]. As a result, the system's machine and grid voltages effectively resemble a $2N + 1$ level inverter with N times the power rating. For details on interleaving and experimental results of a 35 MW drive, the reader is referred to [11].

T. Geyer was with GE Global Research, Munich, Germany; he is currently with the Department of Electrical and Computer Engineering, The University of Auckland, New Zealand; e-mail: t.geyer@ieee.org

S. Schröder is with GE Global Research, Munich, Germany; e-mail: stefan.schroeder@research.ge.com

In both cases, compared to the standard application with one thread, the part count for the N thread system is increased by approximately N times. This impacts on the system reliability and availability in an adverse way if no appropriate countermeasures are taken. On the other hand, the parallel thread arrangement can be exploited to add redundancy [10]. Moreover, in the event of severe faults, the coupling inductors limit the instantaneous currents, confine the fault to one thread, and thus allow the isolation and removal of the faulted thread by an appropriate fault handling strategy.

In more traditional drive systems (with one thread) the reliability can be boosted by adding redundancy within the thread on a device level. This option is typically employed in thyristor based converters (LCIs). In PWM type converters, examples of this approach include redundant semiconductor switches [15] and adding a fourth redundant inverter leg [3]. Alternatively, by installing multiple threads in parallel, redundancy is added on a thread level [12], [17]. Another option is to use multi-phase machines fed by several inverters arranged in parallel, where each inverter feeds a three-phase group of the machine independently [2].

Reliability considerations and an optimized fault handling strategy for shunt-interleaved drive systems are the focus of this paper. Specifically, after outlining the drive system topology and the control scheme in Section II and recalling the required reliability terminology and mathematical formulas in Section III, the paper derives the failure rates of the common part and the threads in Section IV. Section V compares the reliability of a N thread system (without redundancy) with the reliability of a $N + 1$ thread system (with redundancy). Based on the fault classification in Section VI, Section VII proposes a suitable fault handling strategy that is straightforward to implement and is expected to yield an availability close to 100%. Experimental results are provided in Section VIII, and the paper is summarized in Section IX.

II. DRIVE SYSTEM TOPOLOGY AND CONTROL SCHEME

A. Drive System Topology

The modular structure of the drive system topology with its N threads is shown in Fig. 1 [1], [11]. Each thread consists of a grid-side circuit breaker, a transformer, an Active Front End (AFE), which is often also referred to as an Active Rectifier, a dc-link, an inverter (INV), a machine-side circuit breaker and a coupling inductor.

In this paper, the AFEs and the inverters are three-level NPC converters using 4.5 kV and 4 kA IGCTs, see Fig. 2. The machine is a synchronous machine, which is fed by an integrated brushless exciter using a rotating three-phase diode bridge.




Fig. 1: Modular drive system with N shunt-interleaved threads (here $N = 3$). The system consists of the main breaker, N grid-side breakers, N thread transformers, N back to back converters, N machine-side breakers, N coupling inductors and a synchronous machine

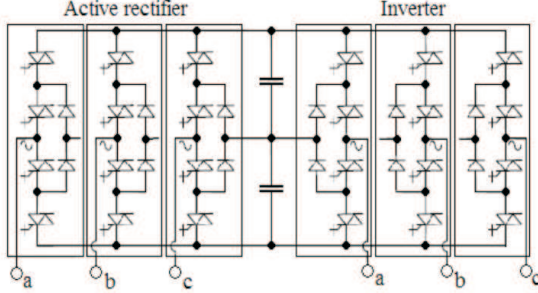


Fig. 2: Detailed schematic of one of the back to back converters, which consists of two three-level NPC voltage source inverters coupled by a dc-link

B. Control Scheme

The controller on the machine side comprises the machine and the N individual thread control processes, which run in parallel. The machine process controls the machine current, which is the instantaneous sum of the N inverter currents. The machine controller is based on a state-of-the-art field oriented control scheme working in a field oriented orthogonal coordinate system. The flux errors are mapped into a flux producing (direct) current reference, while the torque reference is translated into an orthogonal torque producing (quadrature) current reference. In steady state the machine is completely fluxed via the field current supplied by the brushless exciter.

The N INV processes ensure the equal sharing of the machine current among the threads. This is achieved by N individual processes that regulate the difference between the inverter current and the machine current divided by N to zero. The resulting voltage command is sent to the INV I/O module, which generates the firing pulses and senses the INV currents.

The control scheme on the grid side is simpler. The AFEs control their dc-link voltages individually. A master Phase Locked Loop (PLL) keeps the AFE controllers synchronized to the grid and phase shifted among each other.

The voltage and current waveforms of the machine and grid, respectively, are significantly improved by interleaving the N threads, i.e. by phase-shifting the PWM carrier signals. This concept and its benefits are explained in detail in [11] along with experimental results.

III. RELIABILITY CONSIDERATIONS

Throughout the paper, we will pursue a bottom-up approach to determine and evaluate the reliability of the complete drive system. More specifically, the system is broken up into its components and the failure rates of these components are determined based on manufacturer data and/or field experience.




Fig. 3: Markov model of a system with $N + 1$ identical components out of which one component is redundant

The component failure rates are then aggregated to obtain the expected system failure rate and the overall system reliability.

A. Terminology and Definitions

Before proceeding, it is beneficial to recall some basic terminology and definitions related to reliability [4], [9].

- **Reliability:** The reliability R is the characteristic of a system or component expressed by the probability that it will perform a required function under stated conditions for a stated period of time.
- **Failure rate:** The failure rate λ is the frequency with which a system or component fails. Usually, the failure rate is given in Failures In Time (FIT), where 1 FIT is equal to one failure within one billion (10^9) hours.
- **MTBF:** The Mean Time Between Failures (MTBF) of a system or component is the average time between two failures of the said system or component. Assuming a constant failure rate with respect to time (neglecting early failures and aging) and ignoring the time to recover from a failure, the MTBF is the inverse of the failure rate.

$$\text{MTBF} = \frac{1}{\lambda} \quad (1)$$

The MTBF is usually given in hours.

- **MTTR:** The Mean Time To Repair (MTTR) is the average time to repair a failed system or component. It is usually given in hours and is the inverse of the repair rate μ .

$$\text{MTTR} = \frac{1}{\mu} \quad (2)$$

- **Availability:** Availability is the proportion of time a system or component is in a functioning condition and thus available to provide the specified function. The availability is given by

$$A = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} \quad (3)$$

B. System Failure Rate

Consider a system with k components with the failure rates λ_i with $i \in \{1, \dots, k\}$ where the failure probabilities are independent of each other. In a system without redundancy where all k components are required for the system to work, the system failure rate λ_{sys} is the sum of the component failure rates [9].

$$\lambda_{\text{sys}} = \sum_{i=1}^k \lambda_i \quad (4)$$

Next, consider a system with redundancy, i.e. a system with $N + 1$ identical components out of which at least N are required for the whole system to work. As shown in Fig. 3, a




Fig. 4: Time evolution of the probability $P_1(t)$ that the system operates in degraded mode and the inverse of the system failure rate for $N = 3$ and $\lambda_c = 23'300$, which is equal to the thread MTBF derived in the next section. The straight lines refer to an MTTR_c of 7 days, whereas the dash-dotted lines refer to an MTTR_c of 2 days

Markov model with three states can be used to describe this system.

Assume that the components fail with the constant failure rate λ_c and that they are repaired with the repair rate μ_c . Initially, at time zero, let the system be in the state of normal operation (no fault, $N + 1$ components are available). If any of the components fails, the system moves with the rate $(N + 1)\lambda_c$ to the state of degraded mode operation (one fault, N components are available). With the rate μ_c the failed component is repaired and the normal system operation is restored. If, however, another fault occurs before the failed component is repaired, the system moves with the rate $N\lambda_c$ to the failed state. With the rate μ_c both failed components are repaired and the normal system operation is restored. For this, we assume that two failed components can be repaired at the same time without deteriorating the repair rate.

Let $P_x(t)$ denote the probability that the system is in state x at time t , where x denotes the number of failed components. Hence $P_0(t)$ refers to normal operation, $P_1(t)$ to degraded mode operation and $P_2(t)$ to system failure. For the three states the differential equations for the probabilities at time t are

$$\frac{d}{dt}P_0(t) = P_1(t)\mu_c + P_2(t)\mu_c - P_0(t)(N + 1)\lambda_c \quad (5a)$$

$$\frac{d}{dt}P_1(t) = P_0(t)(N + 1)\lambda_c - P_1(t)(N\lambda_c + \mu_c) \quad (5b)$$

$$\frac{d}{dt}P_2(t) = P_1(t)N\lambda_c - P_2(t)\mu_c \quad (5c)$$

with the initial condition $P_0(0) = 1$ and $P_1(0) = P_2(0) = 0$. Moreover,

$$P_0(t) + P_1(t) + P_2(t) = 1 \quad (6)$$

has to hold. The time evolution of the probability $P_1(t)$ is depicted in the upper part of Fig. 4, using the parameters derived in Section IV, while the lower part of Fig. 4 shows the inverse of the system failure rate as a function of time, $\lambda_{\text{sys}}(t) = P_1(t)N\lambda_c$, which is the transition probability from degraded mode operation to the failed state. It can be seen that




Fig. 5: Reliability modeling of the drive system based on the common part (synchronous machine, etc) and the N threads (back to back NPC converters, thread transformers, etc)

steady-state conditions are reached after roughly 1000 h. This motivates the analysis of the Markov model at its equilibrium. Setting the left hand side of (5) to zero and inserting (6) and (5c) into (5b) yields

$$P_1 = \frac{(N + 1)\lambda_c}{(2N + 1)\lambda_c + N(N + 1)\lambda_c^2/\mu_c + \mu_c}, \quad (7)$$

where $P_1 = P_1(t \rightarrow \infty)$. Inserting (7) into the steady-state system failure rate $\lambda_{\text{sys}} = P_1N\lambda_c$ leads to

$$\lambda_{\text{sys}} = \frac{N(N + 1)\lambda_c^2}{(2N + 1)\lambda_c + N(N + 1)\lambda_c^2/\mu_c + \mu_c}. \quad (8)$$

Since typically $\lambda_c \ll \mu_c$ this equation can be simplified to

$$\lambda_{\text{sys}} = \frac{N(N + 1)\lambda_c^2}{(2N + 1)\lambda_c + \mu_c}. \quad (9)$$

The same result is obtained in [7] with slightly different arguments. Further simplifying (8) using $\lambda_c \ll \mu_c$ again yields

$$\lambda_{\text{sys}} \approx N(N + 1)\frac{\lambda_c^2}{\mu_c}. \quad (10)$$

This can be interpreted as a system that is most of the time in the mode of normal operation ($P_0 \approx 1$). The probability P_1 to be in degraded mode is then approximately $(N + 1)\lambda_c/\mu_c$ when ignoring in (5b) the (unlikely) transition to the failed state. Multiplying P_1 with the failure rate $N\lambda_c$ towards the failed state yields the (approximate) system failure rate λ_{sys} given in (10).

IV. RELIABILITY MODELING

Consider the drive system introduced in Section II that comprises a common part (synchronous machine, master control, etc) and N parallel threads. For our reliability calculations the drive is conveniently modeled using the two aggregated blocks common part and thread as shown in Fig. 5. The failure rates of the common part and the threads are λ_{com} and λ_{thr} respectively.

Common part component	Failure rate (FIT)
Synchronous machine including exciter	8300
Motor auxiliaries	1000
Master control incl. auxiliaries	3000
Main circuit breaker: spurious opening	180
2 thread circuit breakers: failure to open (i.e. to isolate the failed thread)	$2 \cdot 220$
Sum ($\text{MTBF}_{\text{com}} = 77'400$ h)	$\lambda_{\text{com}} = 12'920$

TABLE I: Components and failure rates of the common part




Fig. 6: MTBF as a function of the number of threads N for a drive system without redundancy (black) and with one additional redundant thread (grey)

A. Common Part

As listed in Table I the common part includes the synchronous machine with the brushless exciter, the motor auxiliaries, the master control board with its auxiliaries and the spurious opening of the main circuit breaker. The thread circuit breakers are only partially included in the common block as far as they fail to isolate a faulty thread. This implies that for the faulty thread's two circuit breakers only the failure mode 'failure to open' is included. Throughout this analysis, the failure rates given in FIT are typical values taken from the relevant literature [16] and [8], from the manufacturers and from field experience.

B. Threads

The thread components with their failure rates are shown in Table 2. Again, the failure rates are taken from [16] and [8], and from field experience, except for the failure rates of the IGBTs, diodes and snubbers, which stem from [14]. Regarding the thread's two circuit breakers, the failure modes 'spurious opening' and 'failure to close' are included thus complementing the failure mode 'failure to open' included in the common part.

Similar to the circuit breakers and depending on the failure mode, the failure rates of the air-core inductors need to be partitioned among the common part and the threads. Yet, to simplify the exposition, the air-core inductors are neglected.

Thread component	Failure rate (FIT)
24 IGBTs	24 · 250
36 Diodes (anti-parallel and NPC)	36 · 20
12 Snubbers	12 · 100
6 Gate power supplies	6 · 800
2 Voltage sensors	2 · 200
6 Current sensors	6 · 200
12 Dc-link capacitors	12 · 20
2 Circuit breakers: spurious opening	2 · 180
2 Circuit breakers: failure to close	2 · 290
Converter transformer	2400
Thread control incl. auxiliaries	3000
Cooling with redundant pump	2400
Sum (MTBF _{thr} = 42'918h)	$\lambda_{thr} = 23'300$

TABLE II: Components and failure rates of the threads




Fig. 7: Various MTBFs as a function of the thread MTBF for $N = 3$. The straight lines refer to an $MTTR_{thr}$ of 7 days, whereas the dash-dotted lines refer to an $MTTR_{thr}$ of 2 days. The $MTBF_{com}$ is 77'400h. Please note the logarithmic scale on the vertical axis

Since their failure rate is expected to be low, the resulting error is negligible.

In general, in case of doubt, the failure rates of the thread components were underestimated rather than overestimated. This optimistic view of the thread reliability helps to underline that even in the case of very reliable threads, redundancy on the thread level is a necessity.

V. RELIABILITY ANALYSIS

A. System with N Threads (without Redundancy)

For a drive system with N threads and without redundancy every failure shuts down the complete system. Hence, according to (4), the sum of the component failure rates yields the system failure rate $\lambda_{sys} = \lambda_{com} + N\lambda_{thr}$ and its inverse is the system MTBF.

$$MTBF_{sys} = \frac{1}{\lambda_{com} + N\lambda_{thr}} \quad (11)$$

Assuming the previous section's failure rates, the resulting system MTBF of a standard drive with one thread amounts to 27'609h or 3.15 years. This is close to the 31'000h that are reported in [8] for a one thread 8-10 MW drive system. As the number of threads is increased, the system MTBF quickly drops to low values (cf. black bars in Fig. 6). For four threads, for example, the system is expected to fail on average once per year. Such an MTBF is clearly too low for many applications.

B. System with $N + 1$ Threads (including one Redundant Thread)

To boost the reliability, redundancy can be introduced. On a thread level this is achieved by adding one redundant thread to the existing N threads. Using (9) and assuming the thread repair rate to be $\mu_{thr} = 1/MTTR_{thr}$ yields the system MTBF.

$$MTBF_{sys} = \frac{1}{\lambda_{com} + \frac{N(N+1)\lambda_{thr}^2}{(2N+1)\lambda_{thr} + \mu_{thr}}} \quad (12)$$

Figure 6 shows the significant reliability improvement that is achieved by adding an additional thread. Here, we assume

Fault category	Percentage of faults	Examples
Severe thread faults	10%	Device shorts such as IGBTs, diodes, dc-link capacitors, etc.
Non-severe thread faults	70%	Pump failures, thread control failures, trips due to thermal protection, etc.
Common part faults	20%	Master control failures, machine shorts, etc.

TABLE III: Classification of faults in severe, non-severe and common part faults

an $MTTR_{thr}$ of 7 days and use again the failure rates of Section IV. To better understand (12), the impact of redundancy and the influence of a varying thread MTBF, consider $N = 3$ threads and refer to Fig. 7.

Considering first a non-redundant system, one can see that the overall system MTBF (red line with downward pointing triangles) is clearly dominated by the thread MTBF (magenta line with squares). Increasing the thread MTBF significantly increases the system MTBF. However, even for a thread MTBF of 100'000 h, the overall system MTBF is limited to about 23'000 h.

Next, redundancy is added on a thread level. Consider a system that comprises $N + 1$ parallel threads without the common part. In this system, one thread is redundant. As shown in Fig. 7, the additional thread impressively boosts the MTBF (cyan line with diamonds vs. magenta line with squares). This effect is further enhanced by reducing the MTTR from seven days to two days (dash-dotted lines vs. straight lines), and by increasing the thread MTBF.

However, as discussed before, the real system also comprises the common part. Even though the MTBF of the common part (black line) is expected to be higher than the MTBF of an individual thread, it is still comparatively low thus limiting the overall system MTBF of a redundant system (blue line with upward pointing triangles). For the drive system considered here, the system MTBF basically saturates already at modest thread MTBFs of 30'000 h. This saturation is almost independent from the MTTR.

When the MTBF of a system with redundancies reaches the limitation imposed by the common part MTBF, further improvements of the thread MTBF or a reduction of the MTTR cannot further improve the system MTBF. Instead, to enhance the system MTBF, the common part needs to be made more reliable. This is certainly the case here, where the thread MTBF is assessed to be 43'000 h leading to a system MTBF of 71'500 h (with three threads plus one redundant thread). For a 2 + 1 system the system MTBF is slightly higher at 74'300 h.

VI. FAILURE MODES AND EFFECTS ANALYSIS

Before deriving a fault handling strategy, the drive system's failure modes and their associated effects and propagation paths need to be analyzed and understood. This is done in a Failure Mode and Effects Analysis (FMEA) that considers the majority of all possible failure modes of the complete system. For an $N = 3$ thread system, a total of several hundred faults emerges. These faults can be classified in three categories (cf. Table III):

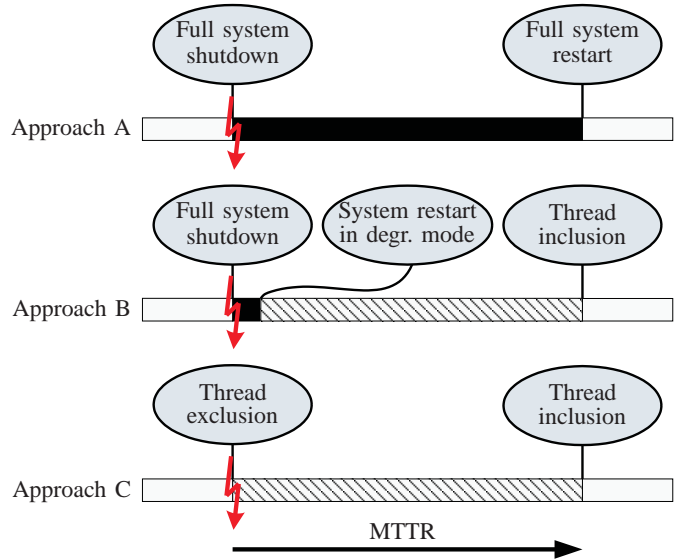


Fig. 8: Summary of the three fault handling approaches A, B and C, and the evolution of the system states in the event of a fault. Light gray refers to the state of normal operation (whole system runs), hatched to the degraded mode operation (system runs, but one redundant thread failed) and black to the failed state (system failed)

1) *Severe thread faults:* Severe (or destructive) faults necessitate adequate and quick corrective actions and fault handling procedures to limit cascading and/or escalating faults and to comply with safety requirements. Severe thread faults include shorts of devices such as IGBTs, diodes, dc-link capacitors, etc. Typically, large currents between the threads, and the machine or the grid are the result, which are only limited by the coupling inductors and the transformer stray inductances. Often, crowbars are triggered to protect the bridge either from over-current or over-voltage. When severe thread faults occur, the thread breakers should be opened as quickly as possible. The FMEA suggests that only about 10% of all faults belong to this class.

2) *Non-severe thread faults:* The non-severe thread faults are less critical. They include pump failures, most thread control failures and trips due to protection (thermal, over-current, under-voltage, etc). They embody the remaining thread faults. Due to the absence of shorts, abnormally high currents do not occur. Therefore, corrective actions do not need to be taken as fast as for severe faults thus providing the fault-handling scheme more time to determine and classify the fault before taking actions. Moreover, a large portion of the non-severe faults does not require any maintenance. These faults include trips due to the various protection measures. Roughly 70% of all faults are non-severe thread faults, out of which some 50% do not require maintenance.

3) *Common part faults:* Faults in the common part of the drive system are by definition single point of failures thus leading to a complete system shut down. These faults could be further partitioned into severe and non-severe faults, which is not required here. Common part faults are master control failures, machine shorts and exciter failures, to name a few. The FMEA indicates that 20% of the faults are single point of failures, out of which the majority is non-severe.

One of the main findings of the FMEA is that the majority of the faults are non-severe thread faults that are relatively

easy to handle.

VII. FAULT HANDLING STRATEGY

A wide range of fault handling approaches are available that provide different degrees of reliability and availability, require different levels of thread redundancy and control software complexity, and vary in terms of their cost and risk. In order to derive a suitable fault handling strategy, the requirements of the process driven by the electrical drive system need to be thoroughly assessed. Based on this assessment and the set of available fault handling approaches, a tailored fault handling strategy with an optimal cost-benefit ratio can be synthesized as will be shown in this section.

A. Fault Handling Approaches

Figure 8 provides a pictorial summary of three viable fault-handling approaches whose characteristics and suitabilities are analyzed hereafter.

1) *Approach A: Traditional Non-Redundant System:* Approach A applies to the standard non-redundant drive system, where, if any component fails, the whole drive system needs to be shut down. To boost the reliability, redundancy on a component level is often introduced, like a second cooling water pump or a redundant thyristor in an LCI stack [16]. Today, the vast majority of drives are installed and operated according to this cost effective solution, which is suitable for applications where a slight reduction in availability is acceptable.

2) *Approach B: Redundant System with Automatic Restart in Degraded Mode:* In many applications, an availability close to 100% is required. Yet, occasional drive system outages can be tolerated provided that the system's outages are short compared to the mechanical and/or thermal inertia of the driven process thus implying that such short outages do not lead to a trip of the whole process. For such an application, a drive system with an automatic restart capability in degraded mode operation is a suitable choice. In case of a fault, the breakers are opened, the complete system is shut down, and the cause and the location of the fault is automatically determined. If the fault is within one thread only, the affected thread is removed from the system by keeping the corresponding breakers open, and the system is restarted in degraded mode with the remaining operational threads. If the electrical machine still spins when the system is restarted, the controller needs to synchronize itself not only to the grid but also to the turning machine [16].

If sufficient redundancy on a thread level is available (i.e. the number of redundant threads is equal to or larger than the number of faulty threads) then full power operation is available also during degraded mode operation. Conversely, the drive is still operational, yet at a reduced power level.

While the system runs in degraded mode, the faulty thread can be repaired or replaced. Once the maintenance has been concluded, this thread can be added on the fly to the running system. We refer to this as thread inclusion. Thread inclusion involves the following steps for the respective thread to be included: Pre-charge the dc-link, close the grid side breaker, start the AFE, close the machine side breaker, start the INV and ramp up the (thread) power to the desired level. On-the-fly

Fault category	Percentage of faults	Fault handling approach	Time of system shut down
Non-severe thread faults	70%	C	Zero, since the system keeps running
Severe thread faults	10%	B	In the range of seconds
Common part faults	20%	A	Some seconds up to the MTTR

TABLE IV: Fault handling strategy for a drive system with parallel threads. The percentage values refer to a three thread system

thread inclusion is an appealing approach that avoids shutting down the complete system when reconnecting a thread.

3) *Approach C: Redundant System with Thread Exclusion:* A few processes do not tolerate even very short drive outages in the range of seconds or below. Such processes include large compressor drives for Liquefied Natural Gas (LNG) trains, which typically require torque glitches to be shorter than one second. This necessitates a sophisticated fault handling strategy. Specifically, an on-the-fly thread exclusion capability is required that avoids shutting down the complete drive system in the event of a thread fault. Special attention needs to be paid to the bump-less torque transfer from normal operation to degraded mode operation - meaning that the torque transient should be as smooth as possible when excluding a thread.

Thread exclusion involves the following steps for the thread to be excluded: Reduce the thread current to zero and ramp up the thread currents of the remaining threads accordingly. If the remaining threads cannot deliver the pre-fault machine current, the machine current needs to be reduced accordingly. Then, stop the INV, stop the AFE, and open the breakers on both ends of the thread to be excluded.

B. Proposed Fault Handling Strategy

Since this paper is mostly concerned with very large power applications in the range of tens of MVA that predominantly feature stringent reliability requirements, approach A is not a suitable choice. A combination of the approaches A and B might be adequate for processes that tolerate short drive outages as discussed above. For more sensitive processes, however, approach B might also lead to a process trip thus effectively resembling approach A. For these processes, a combination of the approaches A and C would be attractive. Yet, addressing all thread faults (including severe faults) by the approach C requires a complex fault-handling scheme, which is costly to develop and qualify, and which carries a certain residual risk of cascading faults and widespread damage to other threads and the common part.

Regardless of the fault-handling strategy, the maximal achievable system reliability is limited by the common part's failure rate. Moreover, as seen in the FMEA, the majority of the thread faults is non severe. Thus, a fault handling strategy that combines all three approaches A, B and C seems to be an attractive choice. With a modest software complexity the majority of the thread faults (the non-severe faults) can be addressed by the approach C, while the remaining few severe thread faults are handled by the option B. The common part faults require the approach A. Since a large fraction of the common part faults requires no maintenance (grid faults, trips due to thermal protection of the machine, etc) the system can

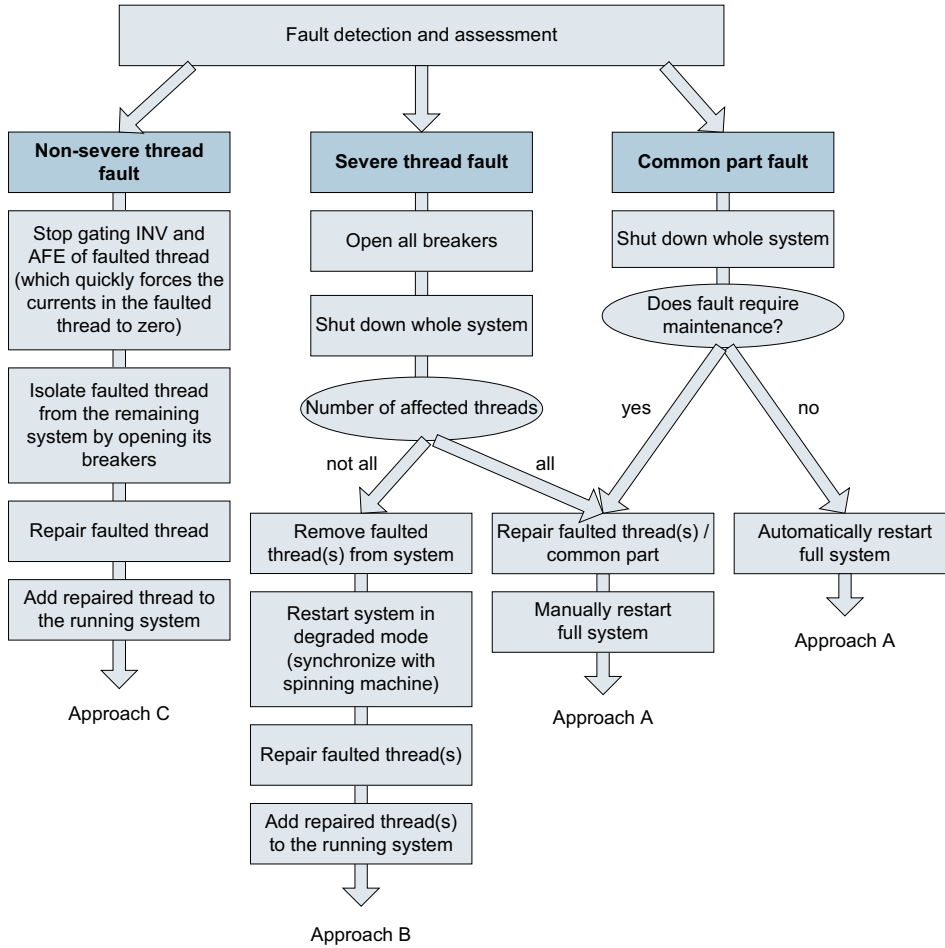


Fig. 9: Detailed flow diagram of the proposed fault handling strategy

be restarted shortly after a trip thus limiting the system shut down time.

From a cost-benefit point of view, this combined A-B-C strategy seems to be the optimum since it maximizes the number of faults addressed by the approach C, minimizes the complexity of the fault-handling scheme, and to a large extent rules out the residual risk of cascading faults. This strategy is summarized in Table IV. A detailed flow diagram is given in Fig. 9.

The availability of the drive system with three threads and with the above A-B-C fault handling strategy in place can be assessed based on (3). Assume that all faults have the same probability. The availability of the drive in the presence of thread faults (non-severe and severe) is effectively 100%. Assume that half of the common part faults have a MTTR of a few seconds and hence an associated availability of 100%, while the other half has the worst-case $MTTR_{com} = 7$ days. Hence the overall system availability A can be approximated by

$$A \approx \left(0.7 + 0.1 + \frac{0.2}{2}\right) \cdot 1 + \frac{0.2}{2} \frac{MTBF_{com}}{MTBF_{com} + MTTR_{com}} \quad (13)$$

yielding 99.98%. This is equivalent to a downtime of 1.75 h per year. Note that the above availability figure includes system operation in degraded mode, where one or more threads are not operational and the available power is reduced accordingly. If full power capability is required also in degraded mode

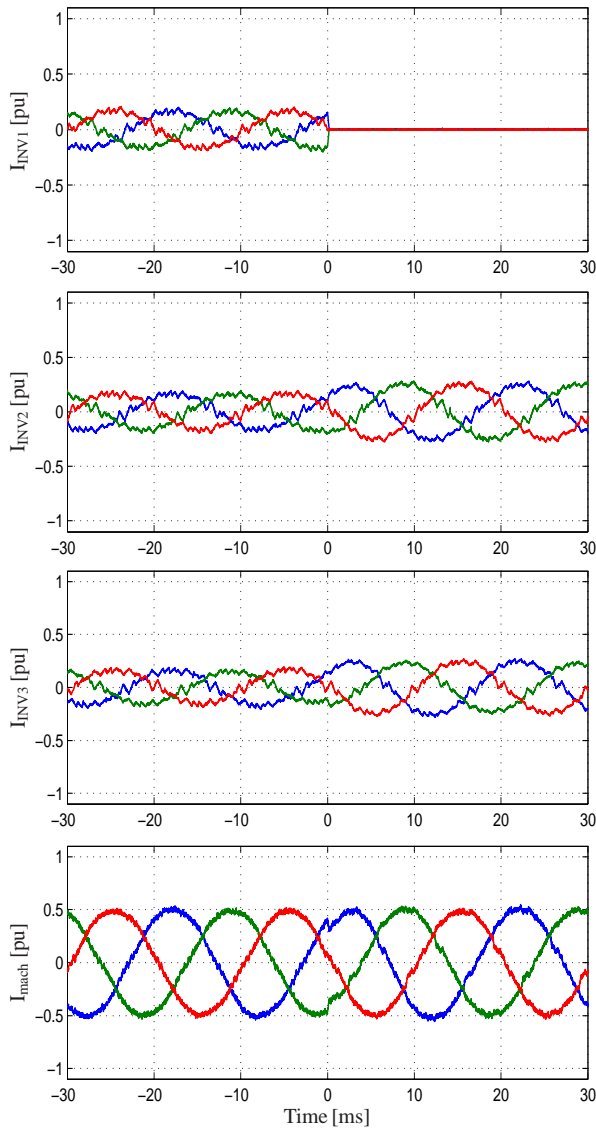
operation, an additional redundant thread needs to be installed. As can be seen from (13), this additional thread does not adversely impact the system availability.

VIII. EXPERIMENTAL RESULTS


In the following we provide selected experimental results. Specifically, thread exclusion during non-severe thread faults was tested on a scaled-down drive system with three threads and a synchronous machine drive rated at 20 kW and 480 V. The threads comprise NPC converters in a back-to-back arrangement. The stray inductances of the thread transformers are sufficiently large to limit the cross currents on the grid side. Accordingly, on the machine side, coupling inductors are used that are of similar value.

We investigated different operating points with different speed and torque settings. In the following we show experimental results at 0.5 pu speed with a torque setting of 0.5 pu. This implies a machine current of 0.5 pu as shown in the lower plot of Fig. 10(a). Each of the three threads provides one third of the 0.5 pu machine current, i.e. 0.167 pu.

Consider a non-severe fault (pump failure or a thermal protection fault, etc) in thread 1 at time zero. The AFE and the INV of thread 1 are stopped within a few control cycles. The thread 1 current drops to zero with a fast transient, which is determined by the coupling inductor. To compensate for the lost thread, the thread current setpoints for the remaining two threads are increased to 0.25 pu. As can be seen in Fig. 10(a),




(a) Fault at 50% load




(b) Fault at 100% load

Fig. 10: INV and machine 3-phase current waveforms of a three thread system during non-severe thread faults at 50% and 100% load, respectively. The AFE and the grid currents are the same. Note that the currents are scaled to their peak values such that for a current of 1 pu its peak values coincide with 1 and -1. These are (post processed) measured waveforms captured on a scaled-down experimental setup



(a) Fault at 50% load



(b) Fault at 100% load

Fig. 11: Torque of a three thread system at 50% and 100% load, respectively, during non-severe thread faults. The torque was reconstructed based on measured machine voltages and currents captured on a scaled-down experimental setup

the thread controllers almost instantaneously ramp up the thread currents from 0.167 pu to 0.25 pu without an overshoot. As a result, the machine remains almost unaffected by the fault. Specifically, the machine currents exhibit only a small perturbation in one of the phases at time zero, while the glitch in the torque (cf. Fig. 11(a)) has a magnitude of less than 20% and is shorter than 5 ms thus resulting in a virtually bumpless transfer. In a last step, the thread breakers are opened to isolate the thread and enable maintenance.

Figures 10(b) and 11(b) show experimental results of a similar thread exclusion scenario at full speed and full load. Since the maximal admissible thread current is only slightly higher than 1/3 pu, the machine current and torque in degraded mode operation need to be reduced accordingly. The apparent unbalance in the currents in Fig. 10(b) is a measurement artefact, which is due to small scaling errors in the current probes used for the measurement system in the test setup. The torque was reconstructed from the measured stator voltages and currents according to the technique presented in [13]. In this reconstruction process, the torque ripple resulting from these measurement artefacts was removed.

Thread inclusion works equally well as experimental results confirm that are not included here due to space limitations.

IX. CONCLUSIONS

Even though modular high power drive systems have a large part count, their modularity can be exploited to provide an overall system reliability and availability that well exceeds the one achieved by standard drives with one thread. To accomplish this, a customized fault handling strategy is mandatory that is preferably easy to implement and maintain, whilst simultaneously avoiding system shut-downs to a large extent.

The majority of faults are non-severe thread faults that can be relatively easily isolated. By excluding and isolating the faulted thread, the overall system can continue its operation in degraded mode. If redundancy on a thread level is available then full power can be provided also in degraded mode operation. Experiments on a scaled-down drive system prove that thread exclusion in the face of non-severe thread faults operates as desired. Specifically, the torque is reduced smoothly (if required) thus limiting the impact on the mechanical load. While the drive system continues its operation, the faulted thread can be repaired and then added on-the-fly to the running system by thread inclusion. As a result, the downtime of such a modular drive system is expected to not exceed a few hours per year.

ACKNOWLEDGMENT

The authors would like to thank Paolo Soldi for the initial simulation results, David Smith and James Nowak for implementing the proposed fault handling scheme, Michael Ozpolat for his help during the experiments and Allen Ritter for his continuous advice.

REFERENCES

[1] R. Baccani, R. Zhang, T. Toma, A. Iuretig, and M. Perna. Electric systems for high power compressor trains in oil and gas applications - system design, validation approach and performance. In *Annual Turbomach. Symp.*, Houston, TX, USA, Sep. 2007.

[2] E. Cengelci, P. N. Enjeti, and J. W. Gray. A new modular motor-modular inverter concept for medium-voltage adjustable-speed-drive systems. *IEEE Trans. Ind. Applicat.*, 36(3):786–796, May 2000.

[3] R. L. de Araujo Ribeiro, C. B. Jacobina, E. R. C. da Silva, and A. M. N. Lima. Fault-tolerant voltage-fed PWM inverter AC motor drive systems. *IEEE Trans. Ind. Electron.*, 51(2):439–446, Apr. 2004.

[4] G. W. A. Dummer, M. H. Tooley, and R. C. Winton. *An elementary guide to reliability*. Butterworth-Heinemann, 1997.

[5] M. Hashii, K. Kousaka, and M. Kaimoto. New approach to a high-power GTO PWM inverter for AC motor drives. *IEEE Trans. Ind. Applicat.*, 23(2):263–269, Mar./Apr. 1987.

[6] J. Holtz, W. Lotzkat, and K.-H. Werner. A high-power multitransistor-inverter uninterruptable power supply system. *IEEE Trans. Power Electron.*, 3(3):278–285, Jul. 1988.

[7] H. H. Kari. *Latent sector faults and reliability of disk arrays*. PhD thesis, Helsinki University of Technology, Espoo, Finland, May 1997.

[8] R.-D. Klug and A. Mertens. Reliability of megawatt drive concepts. In *Proc. IEEE Int. Conf. on Ind. Techn.*, Maribor, Slovenia, Dec. 2003.

[9] K. B. Misra. *Reliability analysis and prediction*. Elsevier, 1992.

[10] U. De Pra, D. Baert, and H. Kuyken. Analysis of the degree of reliability of a redundant modular inverter structure. In *Proc. Int. Telecomm. Energy Conf.*, pages 543–548, Oct. 1998.

[11] S. Schröder, P. Tenca, T. Geyer, P. Soldi, L. Garces, R. Zhang, T. Toma, and P. Bordignon. Modular high-power shunt-interleaved drive system: a realization up to 35 MW for oil & gas applications. In *Proc. IEEE Ind. Appl. Soc. Annual Mtg.*, Edmonton, Canada, Oct. 2008.

[12] B. Shi and G. Venkataramanan. Parallel operation of voltage source inverters with minimal intermodule reactors. In *Proc. IEEE Ind. Appl. Soc. Annual Mtg.*, Oct. 2004.

[13] J. Song-Manguelle, S. Schröder, T. Geyer, G. Ekemb, and J.-M. Nyobe-Yome. Prediction of mechanical shaft failures due to pulsating torques of variable frequency drives. In *Proc. IEEE Energy Conv. Congress and Exp.*, San Jose, CA, USA, Sep. 2009.

[14] P. Steimer, O. Apeldoorn, E. Carroll, and A. Nagel. IGCT technology baseline and future opportunities. In *Proc. IEEE Transm. and Distr. Conf. and Exp.*, Atlanta, GA, USA, Oct./Nov. 2001.

[15] B. A. Welchko, T. A. Lipo, T. M. Jahns, and S. E. Schulz. Fault tolerant three-phase ac motor drive topologies: a comparison of features, cost, and limitations. *IEEE Trans. Power Electron.*, 19(4):1108–1116, Jul. 2004.

[16] P. Wikström, L. A. Terens, and H. Kobi. Reliability, availability and maintainability of high-power variable-speed drive systems. *IEEE Trans. Ind. Applicat.*, 36(1):231–241, Jan./Feb. 2000.

[17] K. Xing, F. C. Lee, D. Borojevic, Y. Zhihong, and S. Mazumder. Interleaved PWM with discontinuous space-vector modulation. *IEEE Trans. Power Electron.*, 14(5):906–917, Sep. 1999.