ON THE RELIABILITY OF SPACECRAFT SWARMS

S. Engelen [(1)], E. Gill [(1)], C. Verhoeven [(1)]

(1) TU Delft, Chair of Space Systems Engineering, Kluyverweg 1, 2629HS, Delft, The Netherlands, +31152781934, s.engelen@tudelft.nl

ABSTRACT

Satellite swarms, consisting of a large number of identical, miniaturized and simple satellites, are claimed to provide an implementation for specific space missions which require high reliability. However, a consistent model of how reliability and availability on mission level is linked to cost- and time-effective design of the individual swarm satellites has not yet been done. We have established a method to model how applied technology and processes for designing swarm satellites under cost and time constraints impact the system-level performance of swarms. The method is applied and discussed for a future astronomy mission using a satellite swarm.
Swarm satellites are severely constrained by mass, as they have to be produced in large numbers. This generally implies that they feature drastically reduced internal redundancy. This is only acceptable when all satellites are functionally identical, and can hence take over certain tasks of a malfunctioning satellite, resulting in a graceful degradation of the system performance. This swarm feature renders it significantly flexible and robust, yet it potentially affects its reliability and system throughput.
Therefore, in this paper we investigated and show how the reliability of an individual satellite transfers into the overall system reliability, and hence the associated throughput of the swarm as a whole. We generated a generic model of a simple swarm satellite, and used it in conjunction with a Markov-chain based reliability analysis to assess the impact of a failure of each of the sub-systems on the functionality of the individual satellite, as well as its impact on the functionality of the swarm. Further analysis was done using Monte-Carlo simulations. The research focussed mainly on the (useful) lifetime of the system as a whole both when considering full and partial failures of elements. This was done using the assumption that swarm elements could still function in a reduced operational state when non-critical components failed. Also, the effect of recoverable malfunctions is investigated.

1 INTRODUCTION

The prospect of satellite swarms in space is slowly gaining acceptance in the space community [1]. The on-going miniaturisation of satellite components, as well as the advent of CubeSats has slowly been enabling the design of cheap yet potentially effective swarm satellites. Two such examples are the OLFAR radio telescope [2], and the QB50 project [1]. The latter attempts to characterise the lower thermosphere by using in-situ measurements taken using a network of 50 satellites.

Whilst swarms of highly miniaturised satellites appear attractive for select scenario's (e.g. OLFAR), physical limits, (e.g. the diffraction limit in optical instruments) will prevent certain systems from being miniaturised enough to allow effective use in a swarm mission scenario.

For those applications where physical limits pose less of a constraint, swarms of highly miniaturised, and thus cheap, satellites can be an attractive design option. This in turn has a severe

impact on the design of the element.
Whilst traditional spacecraft feature internal redundancy and intricate quality and process control for all critical systems [3], these swarm satellites usually have to make do without, because of mass- and cost limitations [4], [5], [6]. Fundamentally, this does not have to be a set-back, as swarm satellites can be regarded as hot-spares of one another, hence can take over functionalities of one another. The impact of a failing satellite on the system however could be significant, as it can result in loss of valuable data.

Also, the lifetime of this type of satellites [7], usually due to their low-cost (non-space qualified) components is limited. Yet this does not necessarily impact the lifetime of the system as a whole.
In this paper therefore, we set out investigating the effects of the design of an individual spacecraft with its subsystems on the overall reliability and lifetime of the swarm as a system.

2 SATELLITE SWARMS AND SWARM SATELLITES

Satellite swarms, following the definition given in [6] consist of multiple functionally identical satellites. This gives them high flexibility, as one satellite can "take over" the role of another, which simplifies or even completely nullifies their dependency on inter-element positioning, in contrast to formation flying satellites. This also has implications on their availability and hence also their required reliability level, which is the topic of this paper. If elements can replace one another, the system can maintain a certain (reduced) level of output even when an individual element is (temporarily) malfunctioning. This is commonly referred to as "graceful degradation" [8], but in this case, it happens on system, rather than on satellite-level.
Swarm satellites are therefore quite independent – in certain cases a single surviving satellite can still produce data, albeit with a severely disrupted schedule, and a severely limited throughput.

3 THE SATELLITE MODELS

In order to assess the lifetime of a swarm and allow comparison with more traditional satellites, a generic satellite model has been defined. This model is based on the functionalities of a satellite, focussing mainly on the data throughput in the system, as well as the interdependence of subcomponents. The model allows assessing system-level performances. In order to allow for a fair comparison, two distinct models are used: one model of a swarm satellite, without internal component redundancy and COTS components, and one of a traditional satellite, featuring internal redundancy and high reliability components.

The swarm satellite model consists of a central On-Board Computer (OBC), which controls all of the element's data flows and operational states. Furthermore, the element consists of an (autonomous) attitude determination and control system (ADCS), a mass storage bank, a radio frequency (RF) inter-satellite link (ISL), a long range RF downlink (DL) and a propulsion unit. The power supply unit (PSU) is considered a single point of failure, as when the power supply system fails, all of the satellite is considered to fail. Mechanical systems such as structures and thermal control subsystems are neglected inasmuch as they do not contribute to the output of the satellite as a whole, and their lifetime is considered irrelevant when compared to the lifetime of active systems.
A payload is also included. This model is based on the current baseline for the requirements for satellites for the OLFAR swarm [2], which attempts to use a swarm of satellites to interferometrically image the radio sky at very low frequencies. Currently, OLFAR requires at least 5 operational satellites in order to allow for interferometry.

In the traditional satellite model, the mean lifetime of all of the internal components has been increased as compared to the lifetimes given in Table 1. No inter-satellite link is considered, which also implies the satellite cannot function as a relay. Furthermore, no analysis has been performed on this satellite's functioning in a swarm, as this is deemed outside of the scope of the paper.

An overview of the swarm satellite model is shown in Figure 1. The traditional satellite model is not shown due to the high degree of similarity.

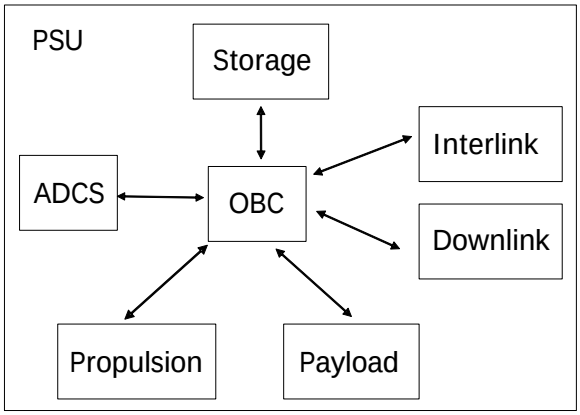

Figure 1: The swarm satellite model. The OBC is considered as the central hub in this satellite, whilst all systems depend on the availability of the power supply unit (PSU).

From this model, it is immediately apparent that a failure of either the power supply or the OBC results in a total satellite failure. The other systems will result in different failure cases, in which the satellite is able to continue operations, for example as a relay station.

4 MEAN-TIME-TO-FAILURE (MTTF) DETERMINATION USING MARKOV CHAIN ANALYSIS

In order to allow analysing the availability, and also reliability of satellites, a mean-time-to-failure analysis was pursued. This analysis renders the operational lifetime of a satellite, which can be used to allow assessing the lifetime of a system of satellites in an analogous manner.
Determining the MTTF is traditionally performed using a Markov Chain analysis [9]. In this process, failure rates of individual components are used to determine the average time for a system to go from an operating state into either a partially failed or even completely disabled state.

When assuming the failure rates of all individual components can be estimated, for example by using their design life time, one can derive a Markov Chain for a certain system. Taking a simple branch of such a chain, as shown in Figure 2, one can represent the possible states of operation with a number (in the example 1, 2 and 3). State 0 represents the nominal state, in which all systems are operational, and 1 represents the state where a single subsystem has failed. States two and three then represent a state where another individual subsystem has failed, after the first subsystem had already failed. Furthermore, λ_1, λ_2 and λ_3 represent the failure rates between the different states.

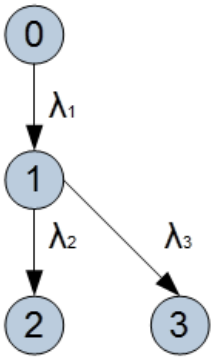

Figure 2: An elementary branch of a Markov Chain

Then this branch can be represented by the set of partial differential equations, as in Eq. (1), (2) and (3):

$$\frac{\partial P_0}{\partial t} = -\lambda_1 P_1 \tag{1}$$

$$\frac{\partial P_1}{\partial t} = \lambda_1 P_0 - (\lambda_2 + \lambda_3) P_1 \tag{2}$$

$$\frac{\partial P_2}{\partial t} = \lambda_2 P_1 \,, \qquad \frac{\partial P_3}{\partial t} = \lambda_3 P_1 \tag{3}$$

In which P_1 represents the chance of being in state 1.

Then, taking $\lim_{t\to\infty} P_1$ will result in the time taken for the system to arrive at P_1. This procedure can be repeated for any branch, and consequently for all possible states. The total system failure time (the MTTF) can then be represented by the addition of all individual failure times.

Our system is a lot more complex as it contains a lot more states, especially when considering recoverable errors, which are treated in more detail later in this section, as well as in section 5.3. The basic principle remains, however. The state-tree for the swarm satellite is shown in Figure 3. The model contains 27 possible states, and for each level, which represents the number of sequentially failed components in the chain, a distinct failure rate per component is used.

Up to three sequential component failures are (in select cases) assumed to allow for continued operation of the element in the swarm. A failed inter-satellite link is also considered in the tree, though this implies the satellite will not be able to function as part of the swarm anymore, and will therefore act autonomously, and independently.

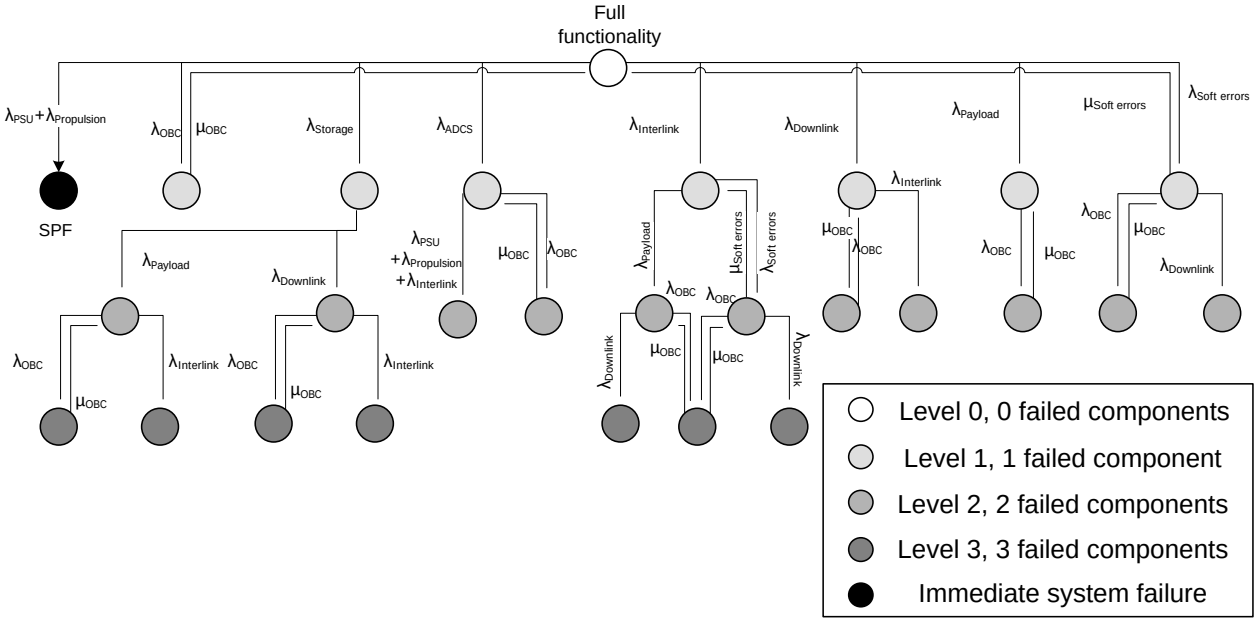

Figure 3: The (simplified) Markov Chain of the satellite model. Darker shades indicate more sequentially failed components.

The satellite structure is not taken into account, and neither is thermal control, as failures are considered to be of an electrical or digital nature. This assumption is based on the notion that in general, structural components outlive electronic systems.

One extra state is considered in the Markov Chain: so-called “Soft errors”. These include single event upsets (SEU/SEL) and general software errors. Their likelihood of occurring is quite high [10], yet they are recoverable [11], for example through introducing internal redundancy (e.g. triple-voting systems), or through fault-acceptance. In case of fault acceptance, the system is allowed to continue with an error up to a certain (predefined) point in time, after which the system is reset to a known state, for example by scrubbing the memory [12] of a certain system.
This category is added specifically to distinguish between a “hard failure”, in which the component is damaged or otherwise incapacitated in a permanent fashion, and “soft failures” which could be recoverable by for example rebooting the system. The results for both a system which performs sporadic reboots and a system which does not are shown in later sections.

Note that the design of the swarm satellite features no internal redundancy. Only soft error mitigation techniques are taken into account, and some measure of fault acceptance is taken into account as well, in that the availability requirement is not considered. In contrast, the traditional satellite model features internal redundancy, yet does not feature an inter-satellite link, as the traditional satellite is assumed to operate alone.

5 RESULTS OF THE MTTF ANALYSIS

In order to assess the lifetime of a satellite, a Markov Chain reliability analysis procedure was performed. The lifetimes of the satellite's individual components (in this case the mean time to failure), as shown in Table 1, were taken to be equal to approximately 5 years and used as input parameters. Certain subsystems, such as the OBC and the ISL and RF Downlink were deemed more fault-prone than others, and for higher levels, the systems were considered "stressed", and therefore more error-prone. In all cases, only the lifetime of the specific component was considered. A separate state-branch containing all so called "soft-errors" was introduced, which allows investigating the effects of recoverable errors on the system lifetime, as well as the physical lifetime of a component. This is done specifically to account for the occurrence of soft errors in the internal components of each spacecraft, without neglecting component lifetime-limiting effects such as total ionising dose (TID).

5.1 Lifetime

The lifetime of a single satellite is determined largely by its components, and their inter-dependency. Our model is straightforward, as avalanche-like effects of failures aren't considered.
When considering a satellite swarm, the total system lifetime depends not only on the lifetime of the individual satellites, but also on the required number of operational satellites. This can easily be understood when considering a single satellite. The lifetime of a single satellite depends largely on the lifetime of its weakest link. When (mass-) producing series of similar satellites, their lifetime is assumed to follow a certain statistical distribution, due to minute deviations in component quality.
In our case, we assumed a Gaussian probability distribution of satellite failure rates, both for simplicity and due to a lack of sufficient statistical data.
Now, when considering a swarm of satellites, the redundancy factor comes into play: Consider a swarm consisting of *n* elements, of which *m* have to be operational, with a normally distributed probability of failure:

- In case n=m, there's no redundancy, and the system's lifetime will be defined by the lifetime of the weakest component present in the system. As can be seen in Figure 4 (a), the system lifetime is found at the left extreme of the Gaussian distribution.

- In case $n>m$, redundancy in the system will shift the probability of system failure towards the right of the distribution (see Figure 4 (b)). In an optimally designed system, $n = \alpha + m$, in which "α" represents the expected number of failed satellites prior to reaching the design life of the satellites. In case of a Gaussian probability distribution, this yields: $\alpha = (\mu - \sigma)n$, where μ is the mathematical expectation, and σ is the standard deviation.

- In case $n>>m$ and $n > \alpha + m$, the system lifetime has been pushed beyond the peak of the Gaussian lifetime distribution, and the system lifetime relies on extreme cases, in which select satellites far surpass their design lifetime. This is shown in Figure 4 (c).

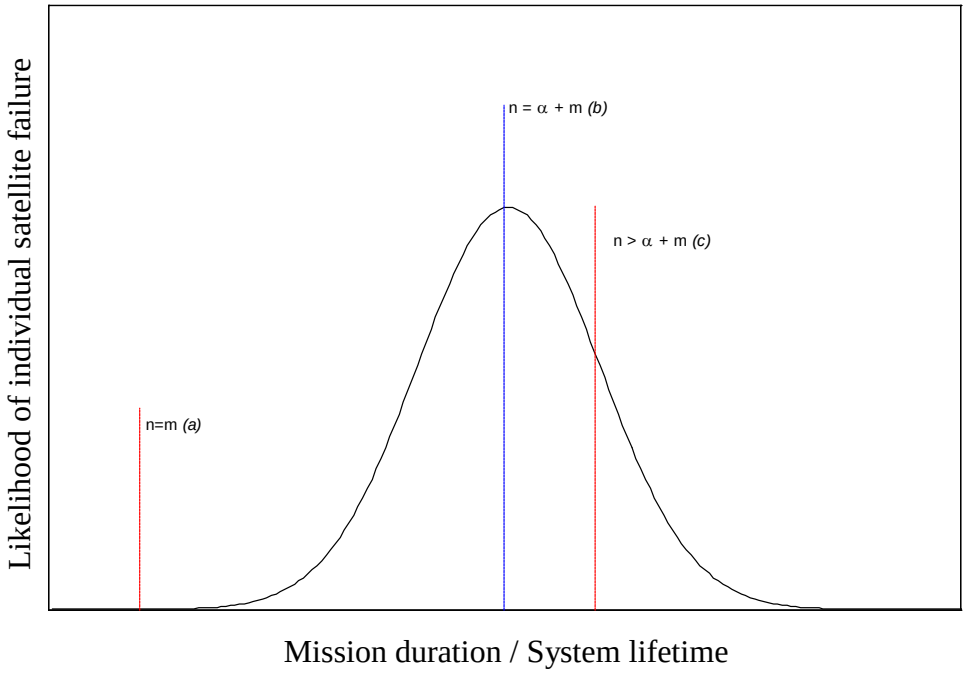

Figure 4: Schematic representation of the effect of redundancy on a satellite swarm, in case of a Gaussian failure probability distribution

The lifetime of the satellite components results in an aggregated satellite lifetime, which is transformed into an overall system lifetime. In case of a single satellite in the swarm scenario, time for the satellite to show a first component failure, without taking soft-errors into account, amounts to 3.4 years, given the design lifetimes of all components for the various stages as listed in Table 1.

The expected lifetime of the traditional (single) satellite amounts to 3.6 years, when taking all component lifetime values out of Table 1 equal to 5 years. Note that both satellites models have components with a design life of approximately 5 years. The traditional satellite model however features internal redundancy for critical components, allowing them to operate for longer. This is modelled by a constant lifetime in all levels of component failure, as no components are expected to operate under stress at any given point in time. The design life of the components hasn't changed however, and also soft errors aren't taken into consideration due to the internal redundancy, as well as the change in design philosophy from fault acceptance to fault tolerance or even fail-safeness.

The distinction between failure rates in different states is taken into account in order to allow considering components under stress. One such scenario could be for example when a satellite has a failed down-link, and another satellite transmits twice as much data as a result of taking over that task from the partially failed satellite. In that case, it is not unlikely that the second satellite's transmitter will give up much sooner due to the higher load.
In case of the traditional satellite, the likelihood of failing with one or more failed components hasn't changed however, as the satellite actually has a reduced workload.

Table 1: Expected component lifetimes for various operational states

Sub-component	Initial expected lifetime	Second level expected lifetime	Third level expected lifetime
	[Years]	[Years]	[Years]
Power supply	5	5	5
OBC	5	1	1
Propulsion	5	1	1
Mass storage	5	1	1
ADCS	5	5	5
Inter-satellite link	5	1	0.1
Downlink	5	1	0.1
Payload	5	5	5
Soft-errors	10 [hours]	10 [hours]	10 [hours]

5.2 Modelling the lifetime of a swarm of satellites

In case of a swarm of satellites, the lifetime can be predicted by considering the swarm as a parallel system of identical systems, as shown in Figure 5, in which *n* represents the total number of satellites, and *m* the number of required operational satellites.

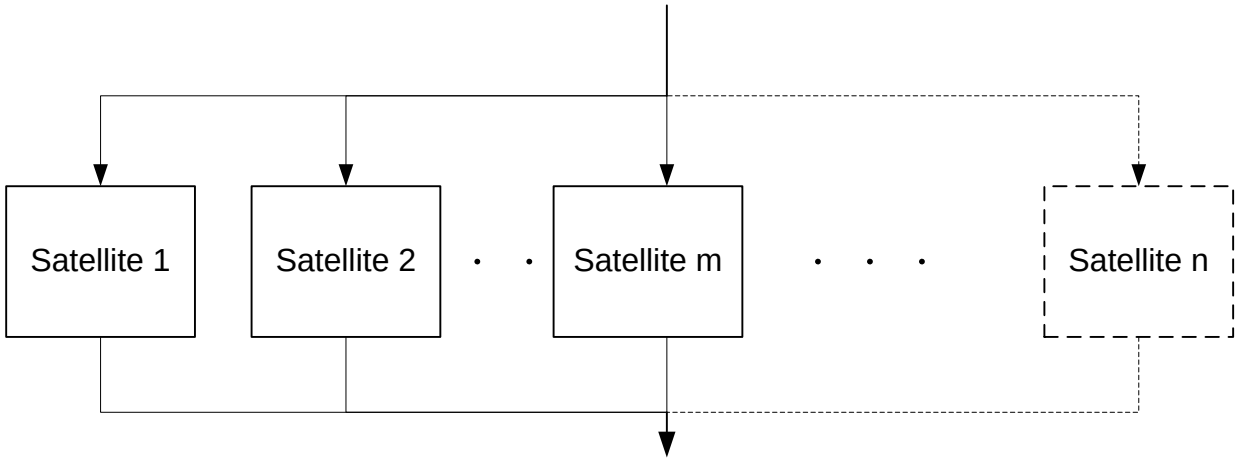

Figure 5: The system model for an *(n-m)* satellite swarm

The lifetime *L* of the system is then definable by Eq. (4):

$$L = \lim_{s \to 0} \left(\frac{1}{s + (n-m)\lambda} \right) \tag{4}$$

in which λ represents the failure rate, or the inverse of the lifetime of a single swarm satellite.

The effect of adding more satellites can be seen in Figure 6. It clearly shows increasing the number of required operational satellites greatly reduced the expected lifetime of the system. Reciprocally, adding extra (spare) satellites into a system will increase the system lifetime.

There is however an optimum, as beyond a certain point, the effect of adding yet another spare satellite diminishes. This is indicated in the figure by the green line, which indicates the maximum effectiveness of adding redundant satellites, by intersecting with the point at which the area under the graph is maximal.
This point coincides with the point at which the system is designed for optimal redundancy, i.e. with $n = \alpha + m$, as was shown in Figure 4 (b).

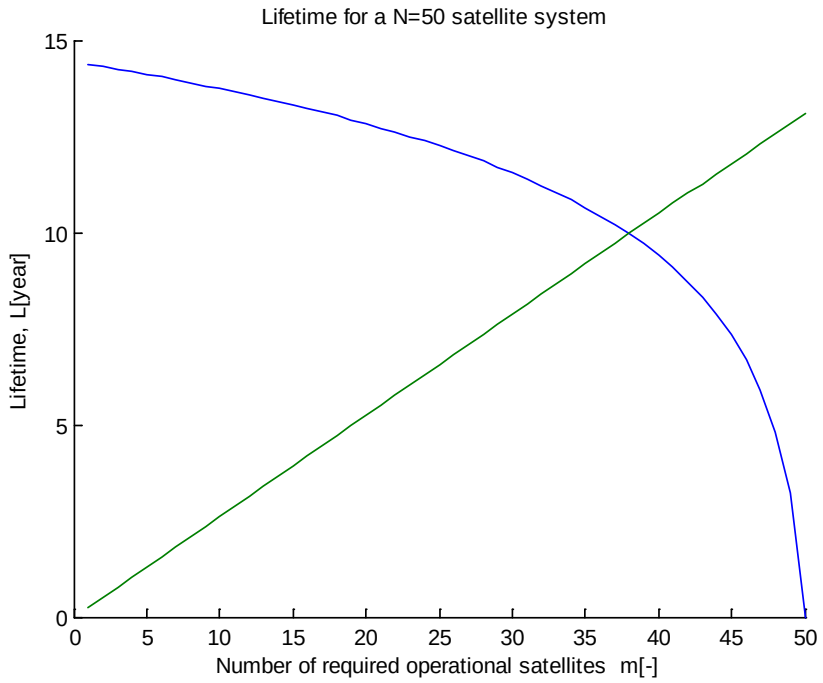

Figure 6: Lifetime prediction for a 50-satellite system.

5.3 Reliability and System Availability

In a complex system, such as a satellite, the system lifetime, as shown in section 5.1, is always lower than their individual component design lifetimes. However, the values computed in section 5.1 still exclude soft-errors, which occur more frequently [10]. When considering soft errors, the lifetime of a swarm satellite reduces from 3.4 years on average to 2.2 years. Those errors can be "repaired" however [11], for example by scrubbing the memory of the component, which gives rise to the question of scrubbing frequency.

The obvious solution to this question is to aim at repairing the problem before it occurs, which implies the scrubbing-rates will vary per system, and they should therefore be equal to or greater than the failure rate of that particular component. However, when the system is being scrubbed, it is probably not operational. This method therefore has implications on the availability of the system, which was the primary reason for scrubbing in the first place. The optimum is therefore dependant on the time taken by the scrubbing process as well. In order to limit the complexity of the analysis, it was assumed the satellite system required a satellite availability of at least 99%, which limits the scrubbing-frequency to 1% of the failure rate per component.

Note that for a system with a high degree of spare satellites, the throughput should not be affected by a single satellite being unavailable for a short amount of time.
The resulting mean-time-to-failures and the corresponding satellite availabilities are shown in Table 2.

Table 2: Computed Mean-Time-To-Failures and satellite availability for various scenarios

Scenario	MTTF per satellite	Satellite availability
Traditional satellite	3.6 years	100%
No soft errors taken into account	3.4 years	100%
Soft errors taken into account, but no scrubbing occurs	2.2 years	100%
Soft errors taken into account, and scrubbing occurs at a rate of 99% of the optimum	3.0 years	99%

5.4 System throughput

Using the values for the mean-time-to-failure, found in Table 1, and the computed values for the satellite MTTF, a Monte Carlo analysis was performed in order to determine the impact of satellites residing in various states of operation during the mission lifetime. The result is shown in Figure 7, with all data normalised to one single satellite unit.
For this analysis, the satellite lifetime was varied between the values given in section 5.3, i.e. satellite lifetimes when excluding soft-errors, when including soft-errors but no recovery procedures, and when using recovery procedures. Also, a normal distribution was applied to the component and satellite lifetimes, with a $\sigma^2$of 0.5.
The radio for the data downlink in each satellite was assumed to have 20% excess available capacity.
For comparison, the hypothetical throughput of the radio downlink and the payload are shown, in the case the satellite would have an infinite lifetime.

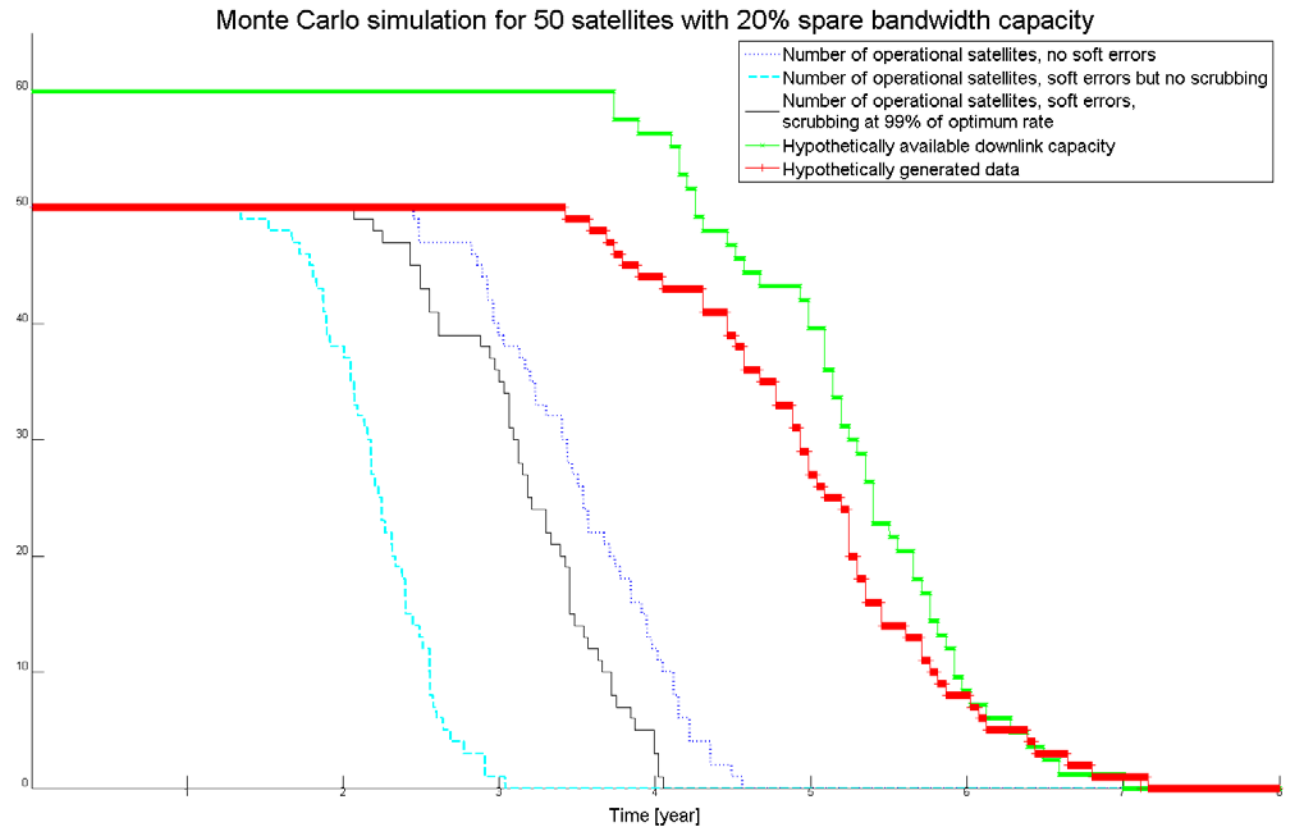

Figure 7: Monte Carlo results for system throughput of a swarm of initially 50 active satellites, with 20% spare downlink capacity, normalised to throughput and lifetime of a single satellite. The satellite and component lifetimes were assumed to have a σ^2 of 0.5.

From the figure, it can be seen that the lifetimes of the satellites are inferior to the lifetimes of the individual components, hence it can be concluded that throughput analysis for swarms can be simplified by only taking full satellite failures into account. Note that the satellite lifetimes as computed in the Markov Chain analysis already included the downlink and payload component lifetimes – the 'hypothetically generated data' and 'hypothetically available downlink capacity' plots - are in fact extracted from those. They do show however, that for the given component lifetime distribution, their lifetimes are of little importance.
Also, the slope of those two are steeper for the downlink than for the payload, due to the higher impact of a loss of a downlink than a payload. As expected, more spare capacity on the downlink is beneficial, as the system will remain operational for a longer time. With the current payload lifetime however and the current excess downlink capacity (assumed to be 20%), it can be seen that this design is close to the optimum, as the downlink capacity becomes insufficient only when merely 9 satellites are remaining.

5.5 Swarm Satellite Design Approach Implications

These results are highly interesting, as due to the discrepancy between the system lifetime and the component lifetime, the need for taking over functionality seems to disappear.
Removing this feature from the swarm satellites, in turn would greatly simplify their design, potentially increasing their inherent reliability.
For a system with parameters similar to the one researched, we have shown that it is justified for a satellite to detect its failure, after which it removes itself from the swarm, rather than to try and continue operations. This has already been suggested in [13], and appears to be a valid approach, as it, rather counter-intuitively, appears to *increase* the reliability of the system as a whole.

6 CONCLUSIONS

We have investigated the effects of component lifetimes and satellite lifetimes on the reliability and availability of satellite swarm systems. We have also analysed the implications on the system throughput.

The results show that the lifetime of a satellite swarm is limited by the lifetime of the individual satellites. This finding, even though the useful lifetime is extendable by adding more satellites, is seemingly analogous to redundancy in monolithic satellites. The cost of adding spare satellites however is also partially compensated for by the additional throughput of the system, making redundancy in a swarm scenario slightly less detrimental to the cost-benefit ratio.

One interesting finding of this research is that the component lifetime, particularly those of the RF downlink or the payload, always exceeds the lifetime of the satellite. The often quoted quality of swarms of being able to take over functionality of malfunctioning satellites has thus little meaning, as the chance of such a situation occurring is minimal. Therefore, a scenario in which satellites simply remove themselves from the swarm when a malfunction occurs is more beneficial than initially thought, and deserves further attention.
Also, it was found that in the case of recoverable errors, their effect could nearly be nullified by proactively repairing their causes, limited only by the effects of the time taken to recover these errors on the system's availability.
The system throughput is affected by failed satellites, yet only due to a loss of their payload functionality. Adding an excess inter-satellite link bandwidth capacity of as little as 20% into the individual satellite design negates the loss of their downlink capability quite effectively.

The above results will allow for a different design philosophy of swarm satellites, in which satellites are added to ascertain throughput, as well as system lifetime, as opposed to increasing satellite reliability through internal redundancy. Also, due to the minor impact of a malfunctioning satellite on the system throughput, the cost of satellites can significantly be reduced e.g. by lowering internal redundancy and employing commercial-off-the-shelf technology.

Swarm satellites should therefore primarily be designed with simplicity as a main design driver – complex modes such as taking over functionality of partially broken satellites do not need to be implemented, as they will effectively reduce the reliability of the system.

7 REFERENCES

[1] J. Muylaert, C. O. Asma en T. E. Magin, „In-Orbit Technology Demonstration Missions for Debris Mitigation," in *4th European Conference for Aerospace Sciences, EUCASS*, St. Petersburg, Russia, 2011.

[2] S. Engelen, C. Verhoeven en M. Bentum, „OLFAR, A RADIO TELESCOPE BASED ON NANO-SATELLITES IN MOON ORBIT," in *24th Anuual AIAA/USU Conference on Small Satellites*, Logan, UT, 2010.

[3] A. Huang, S.-S. Chen, H.-L. Perng en M.-Y. Hsieh, „Reliability allocation and prediction for developing small satellite," in *7th IAA Symposium on Small Sattelites for Earth Observation*, Berlin, Germany, 2009.

[4] C. J. Verhoeven en W. Jongking, „MISAT: A satellite colony," in *8th International Conference on the Commercialization of Micro and Nano Systems, COMS2003*, Amsterdam, The Netherlands, 2003.

[5] C. Verhoeven, M. Bentum, B. Monna, J. Rotteveel en J. Guo, „On the origin of satellite swarms," *Acta Astronautica*, vol. 68, nr. 7-8, pp. 1392-1395, 2011.

[6] S. Engelen, E. Gill en C. Verhoeven, „Systems Engineering Challenges for Satellite Swarms," in *Aerospace Conference*, Big Sky, MT, 2011.

[7] G. F. Dubos, J.-F. Castet en J. H. Saleh, „Statistical reliability analysis of satellites by mass category: Does spacecraft size matter?," *Acta Astronautica*, vol. 67, pp. 584-595, 2010.

[8] W. Nace en P. Koopman, „A Product Family Approach to Graceful Degradation," in *Institute for Software Research*, 2000.

[9] J. Pukite en P. Pukite, Modeling for Reliability Analysis: Markov Modeling for Reliability, Maintainability, Safety, and Supportability Analyses of Complex Systems, 1st Edition, Wiley-IEEE Press, 1998.

[10] C. Dyer, „Radiation Effects on Spacecraft & Aircraft," in *ESA Space Weather Workshop: Looking towards a European Space Weather Programme*, ESTEC, Noordwijk, The Netherlands, 2001.

[11] C. H. Carmichael en P. E. Brinkley, Jr., „Techinques for mitigating, detecting and correcting single event upset effects in systems using sram-based field programmable gate arrays". USA Patent US 7,036,059 B1, 25 April 2006.

[12] R. Naseer, Y. Boulghassoul, M. A. Bajura, J. Sondeen, S. D. Stansberry en J. Draper, „Single-Event Effects Characterization and Soft-Error Mitigation in 90nm Commercial-Density SRAMs," in *Proceedings of the IASTED International Conference*, Kailua-Kona, HI, USA, 2008.

[13] M. G. Hinchey en E. L. Vassev, „METHOD OF IMPROVING SYSTEM PERFORMANCE AND SURVIVABILITY THROUGH SELF-SACRIFICE". US Patent 20100146635, 10 June 2010.