

Swarm Satellites

Design, Characteristics and Applications

Engelen, Steven

DOI

[10.4233/uuid:a5dbc5f7-039b-431b-8f32-d3394e690348](https://doi.org/10.4233/uuid:a5dbc5f7-039b-431b-8f32-d3394e690348)

Publication date

2016

Document Version

Final published version

Citation (APA)

Engelen, S. (2016). *Swarm Satellites: Design, Characteristics and Applications*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:a5dbc5f7-039b-431b-8f32-d3394e690348>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

STEVEN ENGELEN



SWARM SATELLITES

Design,
Characteristics and Applications

Propositions

1. A satellite swarm is a space system consisting of many identical, egalitarian spacecraft, cooperating to achieve a common goal.
2. Monte Carlo analysis should only be used to gain insight into the behaviour of a multivariable system. Due to its stochastic nature, it should not be used to replace in-depth mathematical modelling, however complex that might be.
3. Swarm satellites without propulsion are not going to be able to contribute to the mission goal.
4. The reliability of spacecraft lies far below that of cars, mobile phones or medical devices. The driver behind this discrepancy is mass production, not the engineering methods applied.
5. Nano-satellites are now seen as oddities or even space junk. In the future, people will see monolithic satellites built entirely using dedicated hardware that way.
6. Swarm members are not social entities. Instead, they strive for survival of the species. Any resemblance of a social activity is simply due to the fact this activity increases the chance of survival of the species.
7. Designers are afraid of not fully deterministic processes. Not being that afraid will save quite some energy.
8. Engineers are not scientists. They don't even speak the same language. Methods that have proven to work with scientists will therefore not necessarily work with engineers.
9. We will never encounter *hostile* alien life, until we go out and venture out into the galaxy. If there were any nearby, we wouldn't be discussing it.
10. Nuclear power is safe. It's the powers that be which aren't.

Comments:

- "The powers that be" is an expression, used to indicate the governing power: [https://en.wikipedia.org/wiki/The_powers_that_be_\(phrase\)](https://en.wikipedia.org/wiki/The_powers_that_be_(phrase))

Swarm Satellites

Design, Characteristics and Applications

Swarm Satellites

Design, Characteristics and Applications

Proefschrift

Ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op 3 November 2016 om 10:00 uur

door

Steven ENGELEN

ingenieur lucht- en ruimtevaarttechniek,
geboren te Hasselt, België.

Dit proefschrift is goedgekeurd door de
Promotor: Prof. Dr. E.K.A. Gill
Copromotor: Dr. ir. C.J.M.Verhoeven

Samenstelling promotiecommissie:

Rector Magnificus,
Prof. Dr. E.K.A. Gill
Dr. ir. C.J.M. Verhoeven

Voorzitter
Technische Universiteit Delft, promotor
Technische Universiteit Delft, copromotor

Onafhankelijke leden:

Prof. Dr. ir. P.P. Jonker
Prof. Dr. F.M. Brazier
Prof. Dr. R. Meijer
Prof. Dr. Zhong Hong (George) Zhu

Technische Universiteit Delft
Technische Universiteit Delft
Universiteit van Amsterdam
York University

Overig lid:

Dr. M.J. Bentum

Universiteit Twente



Keywords: Swarm satellites, miniaturisation, reliability, OLFAR, nano-satellites,
distributed space systems
Printed by: Ipskamp Drukkers B.V.
Front & Back: S. Engelen

© Copyright 2016, Steven Engelen

All rights reserved. No part of the material protected by this copyright notice may be reproduced, or utilised in any form or by any means, electronic or otherwise, including but not limited to photocopying, recording or by any information storage and retrieval system, without prior permission by the author.

ISBN/EAN: 978-94-028-0412-6

An electronic copy of this dissertation is available at
<http://repository.tudelft.nl>

*To engineering,
an art form which is all too rarely recognised for its achievements and contributions to society.*

SUMMARY

Satellite swarms at the start of this research were a novelty, and even to date, no practical satellite swarms have proven themselves. In fact, in order to allow in-depth analysis of what would be satellite swarms, a strict definition of what a satellite swarm would entail was required. The definition defined to this end reads “*A satellite swarm can be defined as a space system consisting of many identical, egalitarian spacecraft, cooperating to achieve a common global goal*”, and served as a guideline for all of the analyses performed for this thesis. In order to focus the research into this wide research field, five research questions were formulated and answered during this research. These are:

1. *Which definition of a satellite swarm would be the best fit within the category of existing and planned distributed space system architectures?*
2. *Which types of application areas would be best suited for satellite swarms?*
3. *How to design and optimise a satellite swarm such that it achieves a certain mission goal? What effect does graceful degradation have on this process?*
4. *How to design the swarm elements which, when operated as a satellite swarm, ensure the resulting satellite swarm achieves a given mission goal?*
5. *Which element design would suit the OLFAR mission?*
6. *How to design the most basic swarm satellite?*

The concept of a satellite swarm was and remains extremely appealing, as one can easily imagine that satellite swarms would prove extremely robust and efficient at gathering large volumes of (useful) data. Critical analysis performed during the research has proven that in order to obtain some of the primarily positive aspects of natural swarms such as graceful degradation however, satellite swarms would also have to embrace some of the aspects which are generally considered as negative, for instance using quite drastic methods such as scheduled “suicide”, which prevents defunct satellites from damaging the swarm, as well as (nearly) full autonomy in order to maintain viability of the swarm. This is in stark contrast with tried and true operations for conventional satellites, which are aimed entirely at maintaining recovery options into at least a partially operational satellite. In the case of satellite swarms, this would result in more complex operations and satellite designs, with the resulting increase in total cost, and decrease in satellite reliability. Optimal use of satellite swarms therefore requires a mentality change with both the operators and the designers of satellite systems. This is not that easy to achieve, especially given that the advantages of releasing such constraints primarily show when satellite swarms are operated with larger numbers of elements. Reducing the per-element cost is therefore imperative, as it will either reduce the overall cost of the mission, or at the very least help in reducing the impact of operating less reliable and irrecoverable satellites.

Satellite swarms in turn would significantly increase the throughput of useable data delivered to the operators, as well as enabling missions which would otherwise prove to be economically unfeasible, such as wide-area, low revisit time surveying missions, or even disposable very low orbit in-situ sensing missions, such as the recently proposed QB50-initiative. The niche application areas particularly suited to satellite swarms are those where element interchangeability becomes a benefit, or where the interchangeability does not have an effect on the mission. Satellite swarms, due to the large numbers of spacecraft involved, excel at gathering data on transient, localised or rapidly changing phenomena. However, due to their reduced reliability their availability

may also be somewhat limited, which can result in the swarm missing an event or opportunity in case the swarm's numbers are limited. Given a sufficient number of swarm members and a sufficiently wide field of view per swarm satellite, swarms are very likely to provide at least one observation of a given event, though not necessarily from the optimal observing vantage point.

This requires a different mind-set from mission designers, which have traditionally been able to rely on their satellites unequivocally providing adequate observations. The large numbers also advocate a very high degree of autonomy, as the short inter-satellite distances between the swarm elements provide more adequate and responsive communication links for swarm control purposes, whilst it also allows, or even necessitates offloading much of the micro-management off of the ground segment. In return, using such incentivised, stimergetic or "global goal-seeking" control greatly reduces the communication overhead imposed by transmission of status and control information of each of the elements towards a ground station, which in turn can be re-allocated to payload data. The task of the ground station controllers then moves from computing and verifying control decisions for each of the individual elements to computing, optimising and selecting observational targets and overall swarm configurations through incentives for the entire swarm.

Satellite swarms gain reliability through large numbers. Occurrence of common cause failures however increases with an increasing numbers of satellites, so only non-common-cause failures are avoided through increasing the number of satellites. Due to the large numbers of satellites being launched, the launch cost in a satellite swarm will prove to be one of the larger cost drivers; perhaps even more so than for traditional satellites, as the development cost can be spread across the many members of the swarm. Design changes which have the potential to increase the per-element reliability are to be considered with great care. If a given design change results in a 50% increase in the overall reliability of the element, this will directly translate into a 50% higher reliability of the overall swarm. This in turn could result in a reduction in the number of elements required to guarantee a certain mission duration, which then reduces the launch costs, whilst due to the identical design of a satellite swarm element, the development costs of this design change are only spent once. Also, satellite swarm elements should be designed to at least meet the desired nominal observational lifetime of the overall system, unless replenishments of the swarm are considered. Certain elements will exceed the nominal lifetime, but the drop-off of available elements near the end of the mission could result in a prohibitively high number of satellites to be added, with the associated cost increase.

Degradation of a generalised satellite swarm has been simulated using Markov modelling and subsequent Monte-Carlo analysis of the impact on various properties of the satellite swarm. This has been compared to (a constellation of) traditional monolithic satellites. The traditional satellites benefit from a significantly higher reliability when comparing it to a satellite swarm, which was assumed to be built using Commercial Off The Shelf (COTS) components, similar to those used in nano-satellites today. The cost of a traditional satellite however, and their time-to-flight is significant, which would render the satellite swarm with an economic benefit, as the amount of data returned by the satellite swarm in a given time-span can easily exceed that of the monolithic satellite due to the high number of satellites gathering data. This advantage however is offset when the data to be gathered is not time-critical. Satellite swarms excel at high time-resolution sampling, but due to the reduced reliability, data points will sometimes be lost or may simply never be obtained. Traditional satellites will then remain the system of choice for such missions. As technology advances, also the data gathered by satellite swarms will increase in quality. Simultaneously, traditional satellites will evolve, hence it is not

unlikely that the distinction with satellite swarms will start to blur at some point in the future.

When using elements in a satellite swarm, which have a (proven) reduced reliability compared to their conventional counterparts, simulating the effects on the overall swarm becomes imperative, and should form an integral part of the design process of any such system. This research has shown that it is possible, once a given element design has been defined, to simulate the lifetime, availability and throughput of these elements when operating in a satellite swarm. In turn these simulations allow system designers to appropriately define the number of elements in the swarm in order to achieve the mission goals. As it turns out, for certain mission criteria, current nano-satellite platforms would prove suitable as platforms for a satellite swarm's elements.

The impact of using traditional top-down and bottom-up methods for designing the individual satellite elements has been investigated. Strictly applying either method shows shortcomings, which would result in a high number of iterations in the design process. In an attempt to counteract some of the shortcomings of both methods, an alternative hybrid design approach has been proposed, which relies on behavioural and observability simulations at a very early stage in the design process. This allows the designers to converge more quickly to a suitable design, and limits the amount of over-definition generally required to cover unforeseen issues. No satellite swarm has been flown thus far, or indeed fully designed to date however, which impedes validation of the proposed method. In the STW ASSYS program, a swarm of satellites was being considered for use in a space-borne distributed radio telescope, operating in the low frequency regime (with frequencies of 300 kHz up to 30 MHz). This system is called the Orbiting Low Frequency Array for Radio Astronomy (OLFAR). In this thesis, certain aspects of the OLFAR swarm have been designed using the proposed method, showing the method renders workable results.

Throughout the thesis, the OLFAR swarm has been used as a reference case, and used as an example of what would be an attractive swarm mission. Given that much of the OLFAR mission is still undefined, certain focal points were taken to be analysed. For the reliability analysis, data from past nano-satellite missions was used as a reference. Preliminary analysis shows that a swarm operability of approximately 2.5 years is achievable given the current estimates on the component lifetimes.

A full orbital analysis has also been performed for a lunar science orbit. The lunar science orbits benefit greatly from the irregularities in the Moon's gravity-field. This causes the satellites to naturally drift, which results in the array to be able to fill the U-V-W sphere in which all measurements are projected with new measurements without spending propellant. This drifting period lasts for approximately 100 days, after which the swarm will have to perform a corrective manoeuvre. However the lunar orbit, regardless of the altitude, imposes high relative velocities between the longest baselines. These are currently at least 10 to 40 times above the limits imposed by the desired one second snapshot integration time. Since this integration time of one second is already on the border of what the communication systems would be able to handle, the lunar orbit is currently deemed not yet (technologically) viable. Alternative orbits around the second Lunar Lagrange-point are being studied, showing promising relative velocities. Yet due to the low disturbance forces present in those orbits, many corrective manoeuvres will have to be performed in order to allow for filling of the U-V-W-sphere. High Earth orbits also show promising relative drift rates, but will also demand frequent corrective manoeuvres. In addition, the array will suffer from an increased amount of Radio Frequency

Interference (RFI) generated by Earth, hence the required data transfer rates will increase.

During this thesis, the technical viability of the OLFAR element's power, communication and payload requirements was assessed through prototyping of the relevant systems. The results were more promising than originally thought, indicating that it would be possible to extend three full 10 m tip-to-tip dipole antennas into a 3-unit CubeSat derived platform, as well as deploy low cost solar panels, which would provide up to 30 W of power after having spent up to a year passing through the Van Allen radiation belts, which is currently seen as sufficient, even though for the inter-satellite communications, more power would be beneficial. Those solar panels could include phased array antennas for long distance communication to ground stations on Earth, which increases the directivity, as well as the total antenna area, allowing for a higher data transfer rate.

This research has therefore shown that the methods proposed will allow future swarms, such as the OLFAR telescope, to speed up development, whilst also reducing or at least predicting the technical risks involved in such an endeavour. This in turn could speed up acceptance by mission managers and by extension financiers.

SAMENVATTING

Zwermen van satellieten waren ten tijde van het begin van dit onderzoek een noviteit, en tot op de dag van vandaag is er nog altijd geen satellietzwerm gelanceerd. Een strikte definitie van een wat een “satellietzwerm” inhoudt is daarom een vereiste. Er is dan ook een definitie daartoe opgesteld, en die is gebruikt als leidraad voor alle analyses besproken in deze dissertatie. De definitie die opgesteld werd is: “Een satellietzwerm kan gedefinieerd worden als een systeem dat bestaat uit vele identieke, egalitaire ruimtetuigen, die samenwerken om een gemeenschappelijk doel te dienen”. Ten einde het onderzoek in dit erg brede onderzoeksgebied in te perken zijn er vijf onderzoeksvragen opgesteld en beantwoord gedurende het onderzoek. Deze zijn:

1. *Welke definitie van een satellietzwerm past het beste binnen de categorie van in de ruimtevaart bestaande en geplande gedistribueerde architecturen?*
2. *Welke soorten toepassingsgebieden zouden zich het beste lenen voor het gebruik van satellietzwermen?*
3. *Hoe dient men een satellietzwerm te ontwerpen en optimaliseren opdat de zwerm een bepaald missiedoel bereikt? Welk effect heeft het proces van “graceful degradation” hierop?*
4. *Hoe dient men de individuele satellieten in de zwerm te ontwerpen, opdat ze, wanneer ze gaan samenwerken al seen zwerm, er voor zorgen dat de zwerm zijn doelstellingen haalt?*
5. *Welke satellietontwerp leent zich het beste voor het gebruik in de OLFAR missie?*
6. *Hoe dient men de eenvoudigste zwermsatelliet te ontwerpen?*

Het concept van een satellietzwerm is daarom niet minder aantrekkelijk, voornamelijk omdat het makkelijk is om zich in te beelden, gebaseerd op voorbeelden uit de natuur, dat satellietzwermen zeer robuust zouden zijn, en ook zeer geschikt om grote volumes aan (nuttige) data te verzamelen. Een kritische analyse, uitgevoerd voor deze dissertatie heeft echter aangetoond dat, indien men van een van de grootste voordelen van een satellietzwerm, met name het concept van “graceful degradation”, gebruik wil maken, men drastische maatregelen zal moeten accepteren om dit te bereiken. Deze maatregelen worden normaal gezien als zeer negatief gezien, en houden een soort van geplande “zelfmoord” van de satellieten in, alsook zeer vergaande autonomie. Dit staat in directe oppositie met de gebruikelijke manier van werken bij traditionele satellieten. De manier van werken daar is er volledig op gericht om zo veel mogelijk herstelmogelijkheden in te bouwen om de satelliet zo lang mogelijk in leven te houden. Bij satellietzwermen zou dit leiden tot sterk verhoogde complexiteit op het gebied van operaties en het satellietontwerp, hetgeen daarnaast leidt tot een sterke verhoging van de kosten voor de missie. Ook de betrouwbaarheid van de zwerm zal er onder leiden. Om optimaal gebruik te maken van satellietzwermen zal er daarom een verandering van denken nodig zijn, zowel bij de gebruikers als bij de ontwerpers van dergelijke satellietssystemen. Dit zal echter niet zonder slag- of stoot gebeuren, vooral omdat de voordelen van het loslaten van deze principes zich pas lonen bij het gebruik van voldoende grote aantallen satellieten. Het verminderen van de kostprijs van elke satelliet is daardoor van zeer groot belang, omdat het de kostprijs van de missie verlaagt, of op zijn minst de nadelen van het gebruik van minder betrouwbare, en niet herstelbare satellieten vermindert.

Het gebruik van satellietzwermen daarentegen kan de doorvoer, en de hoeveelheid aan nuttige data die bij de gebruikers van satellietssystemen aangeleverd wordt sterk verhogen, alsmede het mogelijk maken van missies die anders niet haalbaar zouden blijken. Dit zijn bijvoorbeeld missies die grote gebieden simultaan willen observeren, en missies die een

hoge revisitatiefrequentie vereisen. Zelfs wegwerp-satellietmissies, met satellieten die in erg lage banen vliegen om in-situ observaties uit te voeren, zoals het QB50-voorstel, behoren tot de mogelijkheden. De toepassingsgebieden die zich uitermate goed lenen voor het gebruik van satellietzwermen zijn die waar uitwisselbaarheid van de individuele satellieten een voordeel biedt, of waar de uitwisselbaarheid niet uit maakt voor de missie. Door het gebruik van grote aantallen satellieten in een zwerm blinken uit in het verzamelen van vluchtige data, van lokale of snel veranderende fenomenen. Hun beschikbaarheid kan echter ook leiden door hun relatief lage betrouwbaarheid, hetgeen kan leiden tot gemiste kansen tot het observeren van dit soort vluchtige fenomenen indien er niet voldoende satellieten in de zwerm zijn die de observatie zouden kunnen overnemen. Indien er echter voldoende satellieten in de zwerm zijn, en indien deze een voldoende groot bereik hebben is het echter zeer waarschijnlijk dat een zwerm tenminste één observatie van het fenomeen kan maken, zei het misschien niet genomen uit de beste hoek.

Dit vereist echter een verandering in mentaliteit bij de missie-ontwerpers, die traditioneel gezien volledig konden vertrouwen op hun satellieten. De grote aantallen satellieten pleiten ook voor een sterk verhoogde autonomie, vooral gezien de relatief korte afstanden tussen de satellieten onderling kortere responsetijden en daardoor een betere grip op het systeem ten gevolge zouden kunnen hebben. Een welkom neven-effect is dat de gebruikers minder belast worden met het beheren van de dagelijkse, repetitieve taken in de zwerm. Zo een gestimuleerde, stimergetische, of "globaal doelzoekende" manier van besturen van een zwerm vermindert de hoeveelheid communicatie tussen het grondsegment en de zwerm, die daardoor weer vrijkomt voor het gebruik van het doorzenden van nuttige data. De taak van het grondsegment kan daardoor verhuizen van het narekenen en controleren van de individuele beslissingen in de zwerm naar het optimaliseren en uitzoeken van nieuwe observatiedoelen, en daarmee de vorm van de zwerm door het uitrekenen van stimulansen.

Satellietzwermen verkrijgen hun betrouwbaarheid door hun grote aantallen. Indien zich gemeenschappelijke fouten voordoen heeft het verhogen van het aantal satellieten in de zwerm echter niet het gewenste effect, dus het verhogen van de aantallen satellieten heeft enkel een effect op niet-gemeenschappelijke fouten. Doordat er een groot aantal satellieten gelanceerd wordt zal de kosten tot het lanceren ervan een groot deel van de totale kostprijs voor de missie in gaan nemen, misschien zelfs meer dan bij traditionele satellieten. De ontwikkelkosten van de satellieten zelf daarentegen, vermits ze veelal identiek zullen zijn, kunnen over de gehele missie gespreid worden. Ontwerpbeslissingen die mogelijk de betrouwbaarheid van de individuele satellieten zouden kunnen verhogen moet men dan ook zeer kritisch beschouwen. Als de betrouwbaarheid van een enkele satelliet door een bepaalde keuze met 50% verhoogd zou kunnen worden bijvoorbeeld, dan vertaalt zich dit meteen in een verhoging van de betrouwbaarheid van de zwerm met 50%. Dit kan dan op zijn beurt weer leiden tot een vermindering van het aantal satellieten in de zwerm die nodig zijn om de gestelde missieduur te halen, hetgeen op zijn beurt weer kan leiden tot een (drastische) reductie in de lanceerkosten. Doordat het satellietontwerp zelf meestal identiek is voor alle satellieten vertalen de extra kosten voor het verhogen van de betrouwbaarheid van de satellieten zich dan weer niet in een drastische verhoging van de missie-kosten, omdat het extra ontwerp-effort maar één keer gedaan dient te worden. Satellietzwermen moeten overigens zo geschaald worden opdat ze tenminste de gewenste missie-duur halen, tenzij men tussentijdse aanvullingen aanvaardt. Tussentijds aanvullen van het aantal satellieten heeft zo zijn voordelen. Het is een feit dat enkele satellieten waarschijnlijk wel langer dan de gestelde missieduur zullen overleven, maar het zal zeer duur blijken om meteen aan het begin van de missie

voldoende satellieten te lanceren, uitgaande van een klein percentage langlevende satellieten.

De veroudering van een gegeneraliseerde satellietzwerm werd gesimuleerd aan de hand van een Markov model en een daaropvolgende Monte-Carlo analyse van de impact van veroudering op verschillende eigenschappen van de zwerm. Dit is daarna vergeleken met een (constellatie van) traditionele enkelvoudige satellieten. Traditionele satellieten hebben een significant hogere betrouwbaarheid in vergelijking met een satellietzwerm, wanneer er wordt aangenomen dat deze gebouwd is met COTS componenten vergelijkbaar met die die op dit moment gebruikt worden in nano-satellieten. De kostprijs van een traditionele satelliet echter, alsmede hun ontwikkeltijd is hoog, waardoor een zwerm een economisch voordeel zou kunnen hebben, vooral als men kijkt naar het feit dat een zwerm veel meer data naar beneden kan zenden in dezelfde tijdspanne, gezien het grote aantal satellieten in de zwerm. Dit voordeel verdwijnt echter wanneer de data die er verzameld moet worden niet tijd-kritisch is. De zwermen blinken namelijk uit in het snel verzamelen van data, maar door hun verminderde beschikbaarheid kan men er niet altijd van uit gaan dat elk data-punt ook daadwerkelijk verzameld en verstuurd wordt. Voor dat soort missies zullen klassieke satellieten de voorkeur blijven hebben. Geleidelijk aan zal de data die satellietzwermen kunnen verzamelen ook in kwaliteit toenemen. Tegelijkertijd zullen ook klassieke satellieten niet stilstaan in hun ontwikkelingen, dus het is ook niet ondenkbaar dat de scheidingslijn in de toekomst zal vervagen.

Bij het gebruik van elementen in een satellietzwerm die een (aangetoond) verlaagde betrouwbaarheid hebben vergeleken met hun conventionele tegenhangers is het van groot belang om de effecten hiervan op de zwerm als geheel te simuleren. Deze simulaties dienen dan ook een essentieel onderdeel van het ontwerpproces van een zwerm te vormen. Het onderzoek hierin voorgesteld toont aan dat het mogelijk is, eenmaal het ontwerp van een satelliet-zwerm-lid bekend is, om de levensduur, de beschikbaarheid en de datadoorvoercapaciteit van deze elementen in een zwerm te simuleren. Deze simulaties staan de ontwerpers van de zwerm op hun beurt toe om het aantal zwerm-satellieten in een zwerm te definiëren, passend bij de missie-doelstellingen. Zoals uit dit onderzoek blijkt, zouden voor het bereiken bepaalde missies zelfs platformen die lijken op de huidige nano-satellieten kunnen volstaan.

De gevolgen van het gebruik van een traditioneel “top-down” of “bottom-up” ontwerpproces voor het ontwerp van een individuele zwermsatelliet is ook onderzocht. Het strikt toepassen van één van de methoden schiet tekort, hetgeen resulteert in een groot aantal iteraties gedurende het ontwerpproces. In een poging tot het teniet doen van de tekortkomingen van beide methodes wordt een alternatieve, hybride methode voorgesteld, die sterk leunt op observatie- en gedragssimulaties tijdens één van de eerste stappen in het ontwerpproces. Dit staat de ontwerpers toe om sneller tot een oplossing te convergeren, en vermindert het aantal maatregelen dat genomen dient te worden om onvoorziene omstandigheden tegen te gaan. Tot dusver is er nog geen satellietzwerm gelanceerd, of zelfs maar uitontwikkeld, hetgeen validatie van de voorgestelde methode belemmert. In het STW ASSYS programma werd een zwerm satellieten voorgesteld om te gebruiken als radiotelescoop in de ruimte, die zich richt op het frequentiegebied tussen 300 kHz en 30 Mhz. Dit systeem wordt de “Orbiting Low Frequency Array for Radio Astronomy” genoemd, of kortweg OLFAR. Gedurende de thesis zijn sommige sub-aspecten van de OLFAR-zwerm ontworpen volgens de voorgestelde ontwerpmethodes, en de resultaten tonen aan dat de methode werkbare resultaten op kan leveren.

Gedurende de thesis wordt de OLFAR-zwerm als maatstaf gebruikt, en als voorbeeld van een aantrekkelijke zwerm-missie. Gezien het feit dat nog veel van de OLFAR-missie

op dit moment ongedefinieerd blijft zijn enkele focuspunten genomen en in detail geanalyseerd. Voor de betrouwbaarheidsanalyse is dan weer data gebruikt van voorbije nano-satellietmissies. Een voorlopige analyse laat zien dat een zwermlevensduur van ongeveer 2.5 jaar bereikbaar blijkt gegeven de huidige schattingen van de levensduur van de componenten.

Een gedegen analyse van een zwerm in een maanbaan is ook uitgevoerd. De maanbaan geniet voordelen van de onregelmatigheden in het zwaartekrachtsveld van de maan. Dit leidt er toe dat de satellieten onderling uiteen “drijven”, hetgeen de telescoop in staat stelt om de U-V-W-bol, waarop alle metingen worden afgebeeld, te vullen, zonder hiervoor brandstof te moeten gebruiken. Voor de maanbaan is aangetoond dat deze periode ongeveer 100 dagen duurt, waarna de zwerm actie zal moeten ondernemen om weer bij elkaar te komen. De maanbaan zorgt echter, ongeacht de hoogte, voor erg grote onderlinge snelheden tussen de satellieten met de grootste onderlinge afstand. Deze snelheden liggen tenminste 10 tot 40 maal boven de gestelde limieten die volgen uit de eis om een integratietijd van één seconde te handhaven. Aangezien een integratietijd van één seconde al op de grens ligt van wat op dit moment mogelijk lijkt voor het communicatiesysteem wordt de maanbaan voorlopig als niet geschikt beschouwd.

Alternatieve banen rond het tweede Lagrange-punt in het aarde-maan-stelsel werden ook bestudeerd, en de eerste resultaten vertonen beloftevolle onderlinge snelheden. Er zijn in die banen echter zeer weinig verstoringen, waardoor het vullen van de U-V-W-bol veelvuldige manoeuvres zal vergen. Hoge circulaire banen om de aarde zelf tonen ook erg lage onderlinge snelheden, maar ook hier zijn waarschijnlijk frequente baancorrecties nodig. Daarnaast is de interferentie in die banen veel sterker, waardoor de vereiste bandbreedte hoger uit zal vallen.

Gedurende de thesis is ook de technische haalbaarheid van de eisen aan het energievoorzieningssysteem van de OLFAR-satellieten, alsmede het communicatiesysteem en de radio-ontvanger door het ontwikkelen van prototypes. De resultaten hiervan blijken positiever dan oorspronkelijk verwacht, en wijzen er op dat het mogelijk blijkt om drie volle 10-meter lange dipool-antennes uit te vouwen uit een 3-unit CubeSat-achtige satelliet, en tegelijkertijd ook goedkope zonnepanelen, die op hun beurt voor tot 30 Watt aan vermogen zouden kunnen leveren, zelfs nadat ze een jaar lang door de stralings gordels om de aarde hebben gevlogen. 30 W wordt op dit moment gezien als voldoende vermogen, hoewel voor de inter-satellietverbinding altijd meer vermogen zou kunnen gebruiken. In deze zonnepanelen kunnen ook “phased-array antennes” geïntegreerd worden om te gebruiken voor de lange-afstandsradioverbinding met de aarde. Dit levert een smallere bundel op, en een groter vangend oppervlak, hetgeen een grotere doorvoer mogelijk maakt.

Dit onderzoek heeft daarmee laten zien dat de voorgestelde methodes het ontwerp van toekomstige zwermen, zoals de OLFAR telescoop, kunnen versnellen, en gelijktijdig ook het voorspellen van de risico's die het gebruik van een zwerm met zich mee dragen. Op hun beurt kan dit dan weer de acceptatie bij missie-managers versnellen, en daardoor dan ook die van financiers.

PREFACE

Let me start by stating that the OLFAR mission, in the framework of which most of my research has taken place, is truly challenging in all aspects, most of all the technical ones. Will it revolutionise our knowledge of the universe and its initial phases? Of that I have no doubt. Will it discover new phenomena? Quite probably. Will it result in something useful? Well, that depends: the scientists will be able to collect their Nobel prizes, yet the engineers behind all of the technical hurdles will probably remain (mostly) in the shadows. But that is, as Mark Bentum always says, the way of things. And I think most engineers would be comfortable with that.

The real question is though, will it be a swarm of satellites, and will those satellites be built like nano-satellites? Well, if this research has shown anything, it should be possible. It will be even more challenging than using micro-satellites as a platform, yet it will allow using more relaxed design criteria – purely using TRL9 technologies will surely fail to deliver any form of sufficient in-space computing capabilities for the coming ten to thirty years. Will the satellites be reliable enough? I would expect them to, yet there are lessons we can learn from past missions, and also from the traditional satellite world. Excessive internal redundancy will not be the right solution for swarm satellites, nor will custom cable harnesses help in mass-producing the satellites, yet extensive testing (on ground) will yield unexpected results, which can on ground at least be solved quite a lot more easily...

I would like to thank my promotor, Professor Eberhard Gill, for I believe he had the hardest job of all of the people supervising me. He supported me (albeit perhaps at times a little hesitantly, though that was usually probably justified) in all of my seemingly unrelated cursory interests and researches, yet it has allowed me to expand my horizons, and I have used much of it for the completion of this thesis, so in hindsight, perhaps it was not all in vain. I cannot pass by Dr. Chris Verhoeven, who perhaps it the only person I have met who has a wider field of interest as I do, and perhaps having him as a supervisor has amplified some of that (maybe even on both sides of the desk...). Thanks to him, OLFAR has gotten a proper boost towards actually becoming reality. I do at times think back to the days when Chris, Arash, Prem and I were seen as “mutually exchangeable units”, perhaps not unlike a swarm satellite each. I would also like to thank my colleagues at SSE, in particular the PhD students there (Adolfo, Arash, Jing, Rui, Prem and the others) (and at EWI), as you helped make this possible. Thanks guys for the many discussions, and let’s not forget the great times in between the discussions!

The OLFAR PhD team, Raj, Alex, David and I, as well as Mark Bentum (who ought to be professor by now) pushed each other’s work ahead, and I do believe that all of us have delivered something beneficial to the OLFAR mission as a whole. No to mention the army of students (sometimes lovingly referred to a slaves) who toiled endlessly at their thesis’s: Eric, Erwin, Hester, John-John, Kevin, Martin, Matthijs, Teun, Vashisht, Vignes, Yuri and all the others I might have failed to mention. Also thanks to all the people at the various watering holes and offices I’ve shared: Cees-Jeroen, Pooja, Eric, Senad, Akshay, Paul, are just a few of the (perhaps too) many names and faces and memories that pop up, and I know I am forgetting more than half of them already... And Emily, without your unwavering support, none of this would have materialised.

Oh, and perhaps the most challenging question: will OLFAR be launched soon? If I had it my way, definitely!

Steven Engelen, Rijswijk, 2016

CONTENTS

SUMMARY	VI
SAMENVATTING	XI
PREFACE.....	XV
CONTENTS.....	XVII
LIST OF FIGURES	XIX
LIST OF ACRONYMS	XXI
1 INTRODUCTION	1
1.1 A BRIEF HISTORY OF SPACEFLIGHT	1
1.2 SIZE MATTERS!.....	2
1.3 STRENGTH IN NUMBERS?.....	5
1.3.1 Natural swarms.....	5
1.3.2 Satellite swarms	7
1.4 SATELLITE SWARMS AND SWARM SATELLITES	9
1.5 RESEARCH QUESTIONS & METHODOLOGY	10
1.6 THESIS OUTLINE	12
2 SATELLITE SWARMS.....	15
2.1 TYPES OF SATELLITE SWARMS	16
2.1.1 Satellite clouds	17
2.1.2 Foraging satellite swarms.....	18
2.1.3 Schooling satellite swarms	20
2.2 TAXONOMY OF DISTRIBUTED SPACE SYSTEMS	21
2.2.1 Orbital distribution	22
2.2.2 Constituency	23
2.2.3 Control strategy and hierarchy.....	23
2.2.4 Communication methods	28
2.2.5 Synthesis	30
2.3 SWARM MANAGEMENT.....	32
2.4 ENABLING FACTORS	33
2.4.1 Technology: Spin-in.....	33
2.4.2 Platform: Standardisation	34
2.4.3 Architecture: Increased autonomy	38
2.4.4 Paradigm shift.....	40
2.5 SATELLITE SWARM APPLICATIONS.....	42
3 DESIGN OF A SATELLITE SWARM	45
3.1 SIZING OF A SWARM.....	48
3.1.1 The effects of numbers	48
3.1.2 Swarm spacecraft lifetimes and reliabilities	51
3.1.3 Overall swarm lifetime	64
3.1.4 Availability and throughput.....	67
3.1.5 Swarm system reliability.....	72
3.2 REDUNDANCY.....	76

4	DESIGN OF A SATELLITE SWARM ELEMENT	79
4.1	COMMON SATELLITE DESIGN METHODS	79
4.1.1	Top-down swarm satellite design.....	80
4.1.2	Bottom-up systems engineering methods	80
4.1.3	Comparison.....	81
4.2	ALTERNATIVE DESIGN METHOD	83
4.2.1	Element design procedure.....	83
4.2.2	Impact on the swarm design.....	85
4.2.3	Verification.....	86
5	THE OLFAR SWARM	91
5.1	SCIENCE CASE.....	91
5.1.1	Science at low frequencies.....	93
5.1.2	Low frequency radio astronomy	95
5.1.3	Low frequency radio astronomy: From space	98
5.2	OLFAR SWARM IMPLEMENTATION	103
5.2.1	Orbit design	104
5.2.2	Swarm distribution methods.....	117
5.2.3	The OLFAR swarm.....	119
5.3	OLFAR ELEMENT DESIGN.....	121
5.3.1	Overview	121
5.3.2	Payload.....	125
5.3.3	Communication.....	129
5.3.4	Ranging and clock synchronisation	131
5.3.5	Energy supply.....	131
5.3.6	OLFAR element reliability	134
5.4	PREDICTED PERFORMANCE	136
6	CONCLUSIONS	141
6.1	WHAT MAKES A SATELLITE A SWARM SATELLITE?	144
6.2	OUTLOOK	147
	REFERENCES.....	149
	APPENDICES	165
	APPX. A: MEAN TIME TO FAILURE ANALYSIS USING MARKOV MODELLING	167
	APPX. B: LIST OF PUBLICATIONS BY THE AUTHOR.....	169
	APPX. C: ABOUT THE AUTHOR.....	171

LIST OF FIGURES

Fig. 1.1: (a) <i>Meganeuropsis</i> fossil, (b) <i>Sphexomyrma freyi</i> , fossilised in amber	5
Fig. 1.2: Examples of swarming in nature.....	7
Fig. 1.3: Examples of distributed satellite missions.	8
Fig. 2.1: (the lack of) Inter-element position control in a satellite cloud.	18
Fig. 2.2: Control strategy for a harvesting swarm.....	19
Fig. 2.3: Control strategy for a schooling swarm.....	21
Fig. 2.4: Classification tree of distributed space systems (1).....	22
Fig. 2.5: Classification tree of distributed space systems (2).....	23
Fig. 2.6: Classification tree of distributed space systems (3).....	24
Fig. 2.7: Control strategies of distributed space systems.....	26
Fig. 2.8: Communication principles in distributed systems	29
Fig. 2.9: Classification tree of distributed space systems.....	30
Fig. 2.10: Classification tree of distributed systems	31
Fig. 2.11: Number and mass of small satellites launched since 1955.....	35
Fig. 2.12: Various CubeSat compatible satellites.	36
Fig. 2.13: Various CubeSat compatible components.	37
Fig. 3.1: The proposed systems engineering method for satellite swarms	47
Fig. 3.2: Hypothetical probability distribution	49
Fig. 3.3: k-out-of-m model of a satellite swarm	50
Fig. 3.4: Lifetime improvement prediction for an $e = 9$ satellite swarm.....	51
Fig. 3.5: Data-centric model of the swarm satellite as used in the Markov model.....	52
Fig. 3.6: Data-centric model of a decentralised swarm satellite	53
Fig. 3.7: The Markov chain for the centralised swarm satellite model	53
Fig. 3.8: The Markov chain for the decentralised satellite model (a)	55
Fig. 3.9: The Markov chain for the decentralised satellite model (b).....	57
Fig. 3.10: MTTF and MTTF for a simple system	60
Fig. 3.11: Time spent in reduced operational states	61
Fig. 3.12: Time spent in the level 1 reduced operational states	62
Fig. 3.13: Time spent in level 2 of the reduced operational states	63
Fig. 3.14: Time spent in level 3 of the reduced operational states	63
Fig. 3.15: Monte Carlo simulation result for the lifetime of a 100-satellite swarm.....	65
Fig. 3.16: Monte Carlo simulation result for the lifetime	66
Fig. 3.17: Monte Carlo simulation result.....	67
Fig. 3.18: Distribution of the MTTF and MTTF (1).....	68
Fig. 3.19: Distribution of the MTTF and MTTF (2).....	69
Fig. 3.20: MTTF versus availability for the decentralised swarm satellite model.....	70
Fig. 3.21: MTTF versus availability for the centralised swarm satellite model.	70
Fig. 3.22: Effect of failing satellites on the throughput of a 50-satellite swarm.....	71
Fig. 3.23: Reliability of a swarm over time.....	72
Fig. 3.24: Reliability of plot of the centralised swarm satellite	73
Fig. 3.25: Reliability of plot of the decentralised swarm satellite (1)	74
Fig. 3.26: Reliability of plot of the decentralised swarm satellite (2)	74
Fig. 3.27: Reliability of plot of the decentralised swarm satellite (3).....	75
Fig. 4.1: A global overview of the proposed satellite swarm element design method.....	84
Fig. 4.2: The process flow for the proposed design method.....	87

Fig. 4.3: Verification process of a given swarm design	88
Fig. 5.1: Distributions of neutral hydrogen in the galactic plane.....	92
Fig. 5.2: The Radio Sky at 408 MHz.....	93
Fig. 5.3: The cosmic microwave background at 94 GHz	93
Fig. 5.4: The first All-Sky Map, generated by the Planck satellite	93
Fig. 5.5: A brief overview of the expansion of the universe	94
Fig. 5.6: Variations in the Cosmic Microwave Background	95
Fig. 5.7: Geometry and coordinate systems used in synthesis imaging.....	96
Fig. 5.8: NASA’s Explorer 38 satellite.....	98
Fig. 5.9: The effect of occultation by the Moon on signals originating from Earth.	99
Fig. 5.10: The all-sky map at 1.31 MHz, as generated by RAE-2.....	100
Fig. 5.11: Geometries involved in determining the radio eclipse fractions	105
Fig. 5.12: Energy density distribution around the Moon.....	107
Fig. 5.13: Orbital evolution for an initial orbit of 200 km	108
Fig. 5.14: Orbital evolution for an initial orbit of 3000 km	108
Fig. 5.15: Baseline paths in UVW-space	109
Fig. 5.16: Filling of the UVW-voxel sphere for scenario #13.	113
Fig. 5.17: Candidate orbits for the OLFAR array.	115
Fig. 5.18: Operating phases for a given OLFAR satellite	117
Fig. 5.19: OLFAR Element component breakdown	122
Fig. 5.20: FBD for an OLFAR element during the science phase	123
Fig. 5.21: FFD for an OLFAR element during the science phase	124
Fig. 5.22: Overview of the OLFAR satellite geometry.....	125
Fig. 5.23: Image of the SAS system prototype.....	126
Fig. 5.24: Node-level signal acquisition path, with digital band-selection	127
Fig. 5.25: Node-level signal acquisition path, with analog band-selection.....	127
Fig. 5.26: Render of an OLFAR element.....	130
Fig. 5.27: ESS Bus topology as proposed in (Klein, et al., 2013)	133
Fig. 5.28: Prototype of an OLFAR solar panel.....	133
Fig. 5.29: Markov Tree for an OLFAR satellite	135
Fig. 5.30: MTTFF of an n-m swarm consisting of n OLFAR elements	137
Fig. 5.31: MTTF simulation of the number of active baselines over time	138
Fig. 5.32: MTTF simulation of the total accumulated observation time (1)	138
Fig. 5.33: MTTF simulation of the total accumulated observation time (2)	139
Fig. 5.34: MTTF simulation of the reliability of an OLFAR swarm.....	140

LIST OF ACRONYMS

ADC	Analog-to-Digital Converter
ANTS	Autonomous Nano-Technology Swarm
ASIC	Application Specific Integrated Circuit
BMS	Battery Management System
BOL	Beginning Of Life
BPF	Band-Pass Filter
CanX	Canadian Advanced Nanospace eXperiment
CMB	Cosmic Microwave Background
COT'S	Commercial Off The Shelf
CSER	Centre for Satellite Engineering Research
CUTE	Cubical ITTech Engineering Satellite
DARIS	Distributed Antennas for Radio Astronomy in Space
DICE	Dynamic Ionosphere CubeSat Experiment
DS GPS	Deep Space GPS; indicating GPS reception above the orbit of the GPS service vehicles
ECC	Error Control Codes
EOL	End Of Life
ESS	Energy Supply System
FBD	Functional Breakdown Diagram
FFD	Functional Flow Diagram
FFT	Fast Fourier Transform
FOV	Field Of View
FPGA	Field Programmable Gate Array
GEO	Geostationary Earth Orbit
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HEO	High Earth Orbit
ISM	Inter-Stellar Medium
LEO	Low Earth Orbit
LL1, ..., LL5	Lunar Lagrange Point designator
LNA	Low Noise Amplifier
LOFAR	LOw Frequency ARray
MMCX	Micro-Miniature CoaXial
MTTF	Mean Time To Failure
MTTFF	Mean Time To First Failure
OLFAR	Orbiting Low Frequency Antennas for Radio astronomy
O-OREOS	Organism/Organic Exposure to Orbital Stresses
PCB	Printed Circuit Board
PPO/PS	Polyphenylene Oxide/Polystyrene
QB50 project	European 7 th Framework project involving 50 CubeSats
RAE-(n)	Radio Astronomy Explorer n
RF	Radio Frequency
RFI	Radio Frequency Interference
SAS	Science Antenna System
SEU	Single Event Upset
SPF	Single Point of Failure

"Insects have transcended these limits in size, by creating highly social colonies"

*Sir David Attenborough in "Life in
the Undergrowth"*

*"Innovation can't happen without accepting the risk that it might fail. The vast and radical innovations
of the mid-20th century took place in a world that, in retrospect, looks insanely dangerous and unstable."*

Neal Stephenson in "Innovation Starvation"

1 INTRODUCTION

This PhD dissertation concerns satellite swarms, focusing in particular on the design of the satellite swarm's individual spacecraft. It is through the intense cooperation with each other that these satellites form the swarm. Satellite swarms are a new class of distributed space system architectures. The novelty of this architecture brings along notions of inexplicable properties, such as emergent behaviour, which originate from a lack of in-depth understanding of certain aspects associated with the new architecture. One of the goals of this thesis is to describe which properties define a satellite swarm as a distinctly different type of distributed space architecture, and to analyse the benefits and detriments of such a type of architecture. The original idea behind this thesis originated from the MiSat project (Gill, et al., 2007), a large Dutch national research project, which discovered that colonies of low cost, low performance satellites could occupy a niche, in which they would enable missions which otherwise would not be feasible, or at least more expensive to achieve. This idea led to the start of the OLFAR project. OLFAR, which is short for "Orbiting Low Frequency Antennas for Radio Astronomy", is a conceptual study analysing the usefulness of a satellite swarm using low-cost satellites as a platform to form a distributed interferometric low-frequency radio astronomical instrument. To date, OLFAR is one of the more advanced proposals involving a satellite swarm, using risk-management strategies and systems engineering methods similar to the ones considered in this work, and it therefore serves as a reference case for many of the topics treated in this thesis.

This thesis attempts to tackle the problem of designing a satellite swarm's individual spacecraft, which, irrespective of the traditional design problems also has to ensure a swarm formed with such spacecraft performs as intended.

1.1 A BRIEF HISTORY OF SPACEFLIGHT

The launch of Sputnik 1 on October 4, 1957 marked the start of a new era, in which launching artificial satellites was possible. Artificial satellites have since been used for e.g. accurate weather forecasting, telecommunication and navigation purposes, Earth imaging, and astronomy. Sputnik 1 would nowadays classify as a small satellite, in particular a micro-satellite, as it had a total mass of 83.6 kg (see Table 1-1). Sputnik 2, followed one month later, with a launch mass of already 508.3 kg.

The first satellite launched by the United States, called Explorer 1, had a launch mass of 13.97 kg, and Explorer 3, the next successful launch in the US, had a launch mass of 14.1 kg, showing a much more gradual increase in launch mass. Since the field of micro-electronics was still in its infancy at the time, it was not possible to increase the capabilities of those very early satellites without increasing their volume and mass. For many governmental and military satellites, this was indeed the main trend until late in the 1980's, as satellites were outfitted with more potent payloads, as well as a larger collection of payloads per single satellite. As a consequence, only very few small satellites were launched between 1971 and 1989 (see Fig. 2.11), except for the OSCAR radio amateur satellites, which were limited primarily due to financial constraints, and the Russian Strela-1 communication satellites.

This growth in satellite mass had several effects: satellite launch masses increased, which meant ever larger launch vehicles were developed, aimed at decreasing the cost per

kilogram in orbit (for larger launch masses). Small-capacity launch vehicles were discontinued, which in turn increased the cost of launching a small satellite, effectively favouring a larger platform with increased capabilities. The investments required to launch and operate a satellite also grew accordingly, which in turn caused an increase in systems engineering effort and risk consciousness, as well as the advent of specialised insurance companies.

These developments caused a counter-balancing movement, which relied on ride-sharing launches using the spare payload capacity of a launcher when launching a larger satellite. Ride-sharing or piggy-back launching, as this practice is unofficially termed, has not been all that popular, as the main payload of the launcher represents the largest investment. The main satellite could therefore impose restrictions on the ride-sharing satellite, which could limit its usefulness, and therefore also the usefulness of its overall mission and hence the investment. Oftentimes, extensive extra tests were imposed to ensure the safety of the primary payload. A novel highly standardised platform, called CubeSats, which impose a volume of $10 \times 10 \times 10 \text{ cm}^3$ per 'unit', has since taken hold however, which allows bypassing most of these extra tests (see section 2.4.2), which in turn resulted in a significant increase in the number of launched small and very small satellites, starting around 2003.

The large investments involved in traditional, large satellite missions however have caused a shift in mentality as well. In the early days of spaceflight, launch failures and satellite failures were a fact of life, and they were (albeit reluctantly) accepted. With increasing investments however, failures became more expensive, which in turn resulted in an increase in managerial overhead due to the introduction of risk management strategies, as well as a higher engineering effort. This then caused a wave of professionalization in the industry, resulting in more reliable satellites and launchers, yet also less of a pioneering spirit and a significant increase in the cost of space assets. Arguably this has been beneficial in maintaining a low amount of space debris, yet it also dampened creativity and hindered rapid progress. This is especially apparent in the rate of developments: when a satellite launched in the late 1960's or early 1970's had a development time of a few years, satellites in the 1980's and 1990's could easily take 15 years or more to develop, e.g. the Envisat satellite, which took 12 years to develop, and cost upward of two billion euro. In a reaction to these developments, the US Defence Department has therefore established a new office in 2007, called the Operationally Responsive Space Office (ORS), which aims at reducing the development time of military satellites, relying on modularity and standardisation of satellite components and platforms. Other, similar developments are now visible across the globe.

1.2 SIZE MATTERS!

Prior to 1990, satellites with masses lower than a 100-200 kilogram were called "small satellites", and no formal classification existed. That changed, when in 1990 the Centre for Satellite Engineering Research (CSER), at the University of Surrey developed two radio amateur satellites, called UoSAT 1 and 2 (also known by their AMSAT classification OSCAR-9 and OSCAR-11). Those satellites used several microprocessors (Sweeting, 1992) to perform various functions, and were therefore deviating from the then standard means of producing mainly sequencer-based avionics. This gained them the name of "Micro-satellites", after which the decimal mass-based classification strategy,

which has since become commonplace, soon followed (Janson, 2011), (Kramer & Cracknell, 2008). The classification is shown in Table 1-I.

Table 1-I
 SATELLITE MASS CLASSIFICATION CATEGORIES,
 ADAPTED FROM (KRAMER & CRACKNELL, 2008)

Satellite mass category	Mass [kg]
Femto-satellites	< 0.1
Pico-satellite	0.1-1
Nano-satellites	1 - 10
Micro-satellites	10 -100
Mini-satellites	100 - 500

While traditional industrial mechatronic systems are driven by their energy efficiency and operating cost, mass is considered the dominant cost driver for spacecraft. This is primarily due to the cost associated with the launch of a spacecraft, which increases with increasing spacecraft mass. Low cost satellite missions therefore generally aim at reducing the spacecraft mass, up to the point where, for a given functionality, the increased development effort would outweigh the launch cost advantages. At a given point however, satellite masses become so small that a dedicated launch is not economically viable, and such satellites are then launched “piggy-back”, together with a primary, larger mass, satellite. This ride-sharing (Swartvout M. A., 2011) has become extremely popular, and in turn it enabled launching nano-satellites and pico-satellites, (Swartvout M. , 2012). Without ride-sharing, nano-satellite launches would be economically and to date even technologically unfeasible, as no dedicated launchers with such low payload mass capabilities are in use anymore. As satellite designers are forced to reduce the launch mass of the satellites in order to reduce the mission cost, the resulting physical spacecraft size also decreases.

Ride-sharing also limits the amount of available launchers, as well as the choice in terms of target orbits. The primary satellite operators generally limit the ride-sharing customers in terms of volatile and potentially harmful components, as they could pose a threat to the primary satellite. Given the recent boom in nano- and pico-satellite numbers however, it is not unlikely that in the future, dedicated launchers for large clusters and swarms of nano-satellites will either be developed or chartered, due to a lack of a sufficient number of larger satellite launches into the intended target orbits. This also gave rise to the notion of building and operating larger groups of small satellites, benefiting from the effects of larger-volume production; yet at similar launch cost as compared to a single larger monolithic (i.e. constructed as a single piece) satellite with a similar mass, which would later lead to the concept of satellite swarms.

The trend of miniaturisation of satellite components and satellites started in the 1980’s, when the first modern micro-satellites were being built (Kramer & Cracknell, 2008). Miniaturisation implied that traditional components were reduced in size due to a higher degree of integration, generally due to the availability of micro-processors and related components, which allowed reducing the area taken up by control circuitry significantly. In turn, this reduced the required volume taken up by (primarily) the electronics, which reduces the overall mass. Improved electronics also allowed for increases in sensitivity, which in turn reduced the required volume for a given instrument, while also allowing a

higher-degree of on-board (pre-)processing, which reduces the required bandwidth of the spacecraft, which in turn allows using smaller transmitters with smaller antennas, all whilst using less power, which also lead to a reduction in solar panel size. In addition, miniaturised satellites generally reduced the number of payloads aboard a single platform, and in order to save cost, applied, when possible, a common bus.

1

Most nano-satellites and pico-satellites launched or planned to date have adopted the CubeSat standard (California Polytechnic State University, 2013), which was introduced in by Prof. Robert Twiggs and Jordi Puig-Sari in 1999 (Nugent, et al., 2008). The use of the standard forces satellite designers to rethink the concept of a satellite, as it enforces a very high degree of standardisation as well as full-system integration, as both payload and bus components are integrated into a single enclosure. This is different from traditional satellites, which mostly consist of stand-alone payloads and subsystems, joined together by a custom structure and a custom harness. Certain standard platforms, generally referred to as a spacecraft busses exist for traditional satellites as well, yet they are specific to a given manufacturer, and busses can therefore not be interchanged. In the case of CubeSat components, most off-the-shelf components can be interchanged freely. Another prominent feature of nano-satellites is that they almost entirely rely on industrial or commercial electronics, commonly referred to as Commercial Off The Shelf, or COTS components. This results in significant cost savings, as these components were developed for bulk markets, and are therefore produced in very high production volumes. This in turn implies a very low unit cost, even for components applying expensive production processes which have undergone very expensive development cycles. It also ensures that most, if not all, faults in the devices are known quite soon after the release of the component, due to the large number of users involved. One drawback is that components have very short life-cycles, which results in the component only being available and supported by the manufacturer for a few years at most. For traditional satellites, with development cycles of ten years or more, this will be an issue; which can only be solved through buying a large stock of components well in advance. Also, these components were not designed for the space environment. In order to qualify them for application in spacecraft therefore, additional testing and screening is required as well as design practices which circumvent potential issues encountered when operating them in space.

The difference in mentality between miniaturised traditional satellites and nano-satellites applying the CubeSat standard can be seen through a thought-experiment: when scaling up a nano-satellite, none of the platform components truly scale up. Nano-satellites which require larger payloads add a “unit” to their bus, in order to accommodate for the additional volume taken up by the payload, yet all other components remain (more or less) identical. Nano-satellite mission designers encountering a lack of payload capabilities for a given CubeSat size will therefore sooner choose to increase the number of satellites, rather than increasing the size of the payload (although this is also partly due to a lack of availability of large-size CubeSat deployers). In contrast, scaling up a micro-satellite would easily allow for a larger payload, yet the increased dimensions of the satellite also require larger actuators, avionics enclosures, more harness length and therefore more mass, which is allowed since they are not forced to strictly adhere to a standard. In fact, proposals exist which apply a nano-satellite as a whole as a largely self-contained component of a larger satellite, or taken along as a companion satellite. The main driver behind this mentality is cost, as nano-satellites which break with the de-facto CubeSat standard encounter a steep increase in launch- and qualification cost, and no

off-the-shelf payloads are available outside of the CubeSat specification. Traditional (micro-) satellites reach the break-even point between benefiting from adding a second identical satellite much later, as launch costs scale linearly with launch mass, rather than discretely (c.f. a standard CubeSat “unit”); which allows them to benefit from an increase in payload dimensions.

1.3 STRENGTH IN NUMBERS?

Due to the reduced cost of launching single low-mass satellites, it has become possible to launch multiple satellites at cost levels equal to or even lower than a monolithic satellite. Driving this to the extreme, it would be possible to envision large clusters of (low cost) spacecraft in orbit. These clusters could then cooperate to achieve a common goal, which evokes images of insect swarms, flocks of birds or schools of fish, which are deemed to be “more than the sum of the parts” primarily because of the interactions these animals have with each other. Colonies of ants, termites and bees can even be regarded as entirely distributed super-organisms; and it is this property which would be interesting to satellite designers, as ant colonies show advantages over individualistic insects in execution of select tasks.

1.3.1 Natural swarms

In the Paleozoic era (between 541 and 252 million years ago), Earth’s atmosphere was entirely different from today’s atmosphere. The oxygen content was much higher compared to today, with oxygen levels reaching up to 35%. Incidentally, certain groups of insects of that era were considerably larger than their modern day counterparts. A prominent example is a species of dragonfly (*Meganeuropsis permiana*) (Fig. 1.1a), which grew up to a wingspan of about 750 mm, and had an estimated mass of 450 gram. Modern dragonflies in comparison only reach wingspans of up to 190 mm. It has been shown that the increased oxygen content at least partly caused this gigantism (Harrison, et al., 2010), (Verberk & Bilton, 2011), as insects appear to be limited in terms of size due to their limited oxygen-distribution system and the growth rate of their exoskeleton. When the oxygen content reduced in the Triassic era (250-200 million years ago), certain giant insects remained, until they were outcompeted by their smaller counter-parts, or by other animals.



Fig. 1.1: (a) *Meganeuropsis* fossil, (b) *Sphecomyrma freyi*, fossilised in amber (Wilson, et al., 1967)

The earliest evidence for colonial insects dates back to the Mid-Cretaceous era (100 million years ago), with the discovery of an ant species called *Sphecomyrma freyi* (Wilson, et al., 1967)(Fig. 1.1b). The Cretaceous era also saw a rise in oxygen content, up to 25% near the end of the Cretaceous era (65 million years ago), yet that does not seem to have affected ant species, as they had become the dominant insect species by the Middle Eocene (45 million years ago) (Wilson & Hölldobler, 2005). Termites originated sometime during the Cretaceous as well (Thorne, et al., 2000). Bees, which also form highly social colonies, are thought to have originated simultaneously with the first known ants or even earlier (Michener & Grimaldi, 1988). The biomass occupied by a given species on the planet can be used as a measure of success for a certain species. To date, ants (Fig. 1.2 (b)) and termites are arguably the most successful eusocial species of insects on the planet, as they are considered to represent a significant fraction of the total biomass occupied by all insect species.

Most eusocial insect colonies apply a caste system, formed around a central reproductive individual, commonly referred to as the colony's "queen", or in the case of termites one or more male and female reproductive individuals. They contain large numbers of sterile individuals belonging to the worker caste, which collect food and nurture the nymphs and individuals belonging to the soldier caste, in charge of defence activities. Since the colony is lost without its reproductive organs, the colony as a whole can be regarded as a single (distributed) organism, indicating that colonies can be regarded as another way of increasing one's individual biomass.

The number of individuals depends on the species, varying from a few workers up to hundreds of thousands of workers per colony. The largest known super-colony, which is formed when neighbouring colonies coexist peacefully due to their genetic similarities, consisted of an estimated 306 million worker ants, and one million queens (Higashi & Yamauchi, 1979). However, a recent discovery showed that select colonies of Argentine ants appear to form a "mega-colony" spanning the globe, as they appear to be genetically related between different super-colonies (Sunamura, et al., 2009).

A single worker bee's brain contains fewer than a million neurons (Chittka & Niven, 2009). A human brain contains an estimated 85 billion neurons. Typical honeybee colonies however consist of approximately 60,000 worker bees, which would render the bee colony with a total of almost 60 billion neurons, rivalling that of an individual human being (Chittka & Niven, 2009), especially given that larger animals dedicate a larger number of neurons to individual muscle control. This could, to an extent, explain the emergent intelligence seemingly displayed by large swarms of insects, and gave rise to the notion that satellite swarms should also consist of elements containing very little computing power (Verhoeven, et al., 2011).

Insects are not the only animals to form swarms. Social fish and birds show "schooling" or "flocking" behaviour (Fig. 1.2 (a) and (d)); which entails a large number of individuals swimming or flying in close proximity for protection or aero- or hydro-dynamic purposes. Such schools and flocks can also easily contain thousands of individuals. The goal of such a swarm, contrary to most insect swarms and colonies is different however, in that insect colonies focus primarily on foraging for nutrients as effectively as possible, over as large an area as possible. Schooling or flocking is generally a temporary phenomenon, to benefit the overall community of individuals, either in the form of protection or for conservation of energy. Locust swarms (Fig. 1.2 (c)) are different in this respect, in that they swarm in order to harvest food, generally consuming vast quantities of plant-matter along their path.




Fig. 1.2: Examples of swarming in nature
 (a) School of fish, (b) flock of birds, (c) Locust swarm, (d) swarm of safari ants

1.3.2 Satellite swarms

Distributing space systems is traditionally caused by either a demand for larger baselines (e.g. in the case of synthetic apertures), a demand for multi-point (synchronous) sensing or due to a demand for signal fusion (Clement & Barrett, 2002). Initially, global simultaneous coverage resulted in constellations of satellites, such as the Global Position System (Hegarty & Chatre, 2008) or Europe's Galileo constellation and Iridium constellations (Maine, et al., 1995). Experience gained in in-space docking manoeuvres, and the need for precise and very long inter-satellite baselines lead to the appearance of formation flying missions, such as the PRISMA (Gill, et al., 2007) and Grace missions (Kirschner, et al., 2001). Driven further, studies such as NASA's Terrestrial Planet Finder (Beichman, et al., 1999) and ESA's Darwin mission (Rabbia, 2004) both aiming at performing optical interferometry, which requires almost nanometre accuracy, and hence as high a platform stability and positioning accuracy as imaginable.

Coordination of the activities of each of the satellites in these distributed systems requires a complex system for coordination when applying traditional methods of satellite control. Using a control method similar to the local, decentralised coordination and high degree of autonomy and autonomicity¹ (Hinchey, et al., 2005) present in many natural swarms would reduce the operational cost significantly, and indeed allow for complex operations beyond the 40-minute round-trip limit which is currently imposed for near-real-time ground-based control for deep-space spacecraft (Hinchey, et al., 2005).

¹ Autonomic refers to actions taken which are involuntary, such as reflexes




Fig. 1.3: Examples of distributed satellite missions.

(a) Europe's Galileo constellation (Image credit: ESA), (b) ESA's Darwin mission (Image credit: ESA), (c) NASA's GRACE mission (Image credit: NASA), (d) ESA's ATV docking with the International Space Station (Image credit: ESA)

A satellite swarm would share as many useable traits of natural swarms, such as large numbers of *expendable* units, gaining system reliability through shear distribution and per-unit simplicity. Very few mission proposals however, or even concept studies on satellite swarms are (publicly) available. Two prominent ones are the NASA ANTS (Autonomous Nano Technology Swarm) concept mission, and the OLFAR (Orbiting Low Frequency Antennas for Radio astronomy) radio telescope. Under the NASA ANTS concept missions, three different concepts can be distinguished (Hinchey, et al., 2005):

- 1) SARA: The Saturn Autonomous Ring Array, which aims at studying the composition of Saturn's ring in-situ; which consists of a swarm of 1000 pico-class spacecraft
- 2) PAM: Prospecting Asteroid Mission, aims at launching 1000 pico-class spacecraft for exploring the asteroid belt, though they expect to lose 60-70% of these satellites during the transfer phase already.
- 3) LARA: ANTS Application Lunar Base. This concept mission does not form a satellite swarm, but rather uses technologies developed for satellite swarms for a swarm of lunar rovers.

The OLFAR radio telescope (Bentum, et al., 2009), (Engelen, et al., 2010) is intended to study the low frequency regime of radio signals of extra-terrestrial origin, and would consist of an operational array of more than 10 satellites. The mass and final number of OLFAR's satellites is not strictly defined. Cost-wise however, a low mass is highly desirable, as it would allow launching more antennas at a similar cost level, which in turn would reduce the required observation time and increase the scientific output of the

mission, due to the additional baselines. OLFAR’s payload, however, is limited by the throughput offered by the inter-satellite link, which effectively limits OLFAR’s cluster size. Even so, multiple clusters could be launched to obtain multiple observations, in different sub-bands.

Given that the most effective satellite swarms would consist of a large number of elements, a certain amount of autonomy seems desirable, at least from a control perspective. While not strictly required, a large degree of internal autonomy eliminates a high bandwidth-requirement between the swarm and a ground-station, and allows for much tighter formation control, due to the significantly decreased round-trip times for control signals. (Hinchey, et al., 2005) even propose that all swarms should consist entirely out of autonomic elements, as most natural swarms rely entirely on pre-programmed reflexes in each of the individuals. Given the conservative nature of the space industry this is not likely to happen soon, even though most deep space satellites launched until the mid-1970’s used sequencers to control their actions; which in essence are nothing more than rudimentary autonomic systems.

1.4 SATELLITE SWARMS AND SWARM SATELLITES

Satellite swarms are a novelty. This shows in the amount of research being conducted, as well as the lack of actual missions: no mission involving a satellite swarm has been launched to date. Confusingly, ESA recently launched a mission called “Swarm”, yet it covers a constellation of three satellites which do not form a satellite swarm.

Satellite swarms are studied in detail from a control perspective ((Vos, et al., 2013), (Izzo & Pettazzi, 2007), (Morgan, et al., 2013), (Pincirolì, et al., 2008)). The interest lies in controlling a large group of “agents” with as little resource use as possible. This is generally inspired by natural swarms, which show similar traits. Local controllers, which rely almost exclusively on nearest-neighbour communication, are seen as highly beneficial in reducing the complexity of operating a satellite swarm (Jähnichen, et al., 2008).

Design methods and strategies tailored to satellite swarms are explored (Benjamin & Paté-Cornell, 2004), (Winfield, et al., 2005), as well as the reliability of robotic swarms and the overall properties of satellite swarms (e.g. (Bonnet & Tessier, 2007), (Hinchey, et al., 2005), (Levi & Kernbach, 2010)). The reliability of other distributed space systems has been treated in detail, and can be used as a reference (Castet & Saleh, 2012), yet on the specific sub-topic of satellite swarms no detailed reliability analysis has been performed. This is in part attributable to the lack of existing satellite swarms which could be used for validation of the models used in the analysis. Furthermore, the exact definition of a satellite swarm is still an unresolved issue.

The individual swarm satellites, which form the swarm, have rarely been studied in detail. NASA’s ANTS mission proposal has studied the means of locomotion and basic properties of the individual spacecraft (Truszkowski, et al., 2004), and select properties of the OLFAR satellites have been defined (e.g. (Budianu, et al., 2014), (Quillien, et al., 2013)). The design process for a swarm element has only been studied for ground-based robotic swarms (Rutishauser, et al., 2009), (Şahin, 2005), (Rubenstein, et al., 2012), which are used mainly as research platforms for validation of novel control strategies for robotic and natural swarms.

The reliability of swarm satellites, which are still undefined, is therefore unknown, though interpolation from nano-satellite reliability figures could be considered (Monas, et al., 2012) as a starting-point, as the majority of satellite swarm proposals aim at using lower-cost satellites in an attempt to reducing the overall mission cost ((Jähnichen, et al., 2008), (Truszkowski, et al., 2004), (Verhoeven, et al., 2011), (Engelen, et al., 2010)).

As no practical satellite swarms have been flown to date, and only a few mission proposals are publicly available, the OLFAR mission will serve as a test-case used throughout the thesis, and perhaps even as a role-model for future swarm missions.

Radio astronomy from space in contrast is an already established field, with examples like the Planck millimetre-wave telescope (ESA, 2010) observing very short wavelength waves from its vantage point in the second Lagrange point of the Earth-Sun system, and the RadioAstron mission performing VLBI measurements from a Very High Earth Orbit (Kardashev, et al., 2013). In fact, the Explorer 38 satellite, which is also known as the “Radio Astronomy Explorer 1” was already launched in 1968, and it confirmed the existence of Auroral Kilometric Radiation (AKR) originating from Earth’s ionosphere (Weber, et al., 1971). Interferometry, which involves synchronously observing given targets with a long base-line between the observing instruments, has only been done using a single space asset (most notably the RadioAstron Spektr-R spacecraft) and Earth-based radio telescopes. Interferometry in space has been proposed before (e.g. NASA’s TPF (Beichman, et al., 1999) or ESA’s Darwin formation flying satellites (Rabbia, 2004)), yet interestingly only for optical wavelengths, which are the most difficult wavelengths to align, as they are extremely short. Radio-wavelengths have not required any space-based missions, as Earth’s atmosphere is sufficiently transparent, and interferometers on Earth are much easier to construct and operate. For the specific case of long-wavelength radio astronomy however, Earth’s atmosphere is largely opaque, justifying space-based observatories. Given the length of these wavelengths, alignment of the instruments, as well as timing-requirements become much more lenient, which is why satellite swarms could be a worthwhile architecture for such missions (Bentum, et al., 2009).

1.5 RESEARCH QUESTIONS & METHODOLOGY

Satellite swarms attract attention for various reasons. One of the more prominent reasons is the promise of graceful degradation, which is sometimes even touted as an increase in system reliability. While this last aspect can easily be retorted, given that this property does not increase the system’s reliability, but rather allows continued operations in a degraded performance-mode, it could still be a useful aspect to some users.

However, for satellite swarms, and indeed for distributed space systems in general, one can wonder what the long-term effects of graceful degradation are, and under which circumstances they would prove truly beneficial. Also, does such a property emerge for all distributed systems, or are scenarios imaginable in which the system fails before graceful degradation has had the chance to manifest itself?

Furthermore, satellite swarms consisting of hundreds or thousands of spacecraft will not be low-cost missions, if only due to the launch costs involved. Many swarm proposals therefore propose the use of nano- and pico-satellite platforms, which can benefit from COTS developments, and mass-produced components with their origin in related bulk-markets. While this, coupled with for the space industry uncharacteristically large production volumes, would significantly reduce the per-unit cost of such spacecraft, launch cost of a swarm will remain similar to the launch cost of a monolithic satellite of

equal mass. The launch costs for a satellite swarm are arguably even higher, as each of the individual spacecraft will require some form of hold-down and release mechanism during launch.

A properly designed satellite swarm however is indifferent to the effects of the loss of a single spacecraft, which is a property that in turn can be used to replenish or update the swarm with new spacecraft over time (Verhoeven, et al., 2011). This could be used to reduce the launch costs through ride-sharing launches with other satellites, after which the launched spacecraft can join the swarm's operations. The question whether this is economically beneficial will be left to mission planners for each specific mission however, as it depends on a variety of external aspects not inherent to the swarm or the swarm spacecraft. Economic viability is therefore considered to lie well outside of the scope of this thesis.

Whether or not nano- and pico-satellites would present a useful platform for satellite swarms is however a relevant question, especially given their reputation of being unreliable. Their properties of reduced unit cost could still prove beneficial to the overall swarm, provided their reduced reliability can be counteracted by other means, for instance through redundancy or through establishing the availability of graceful degradation. Alternatively, satellite swarms could increase their throughput up to the point where reliability becomes moot, such as missions where all relevant data could be gathered in a much reduced time-frame, during which the reliability of the individual spacecraft could be considered to be sufficiently high. As an example, one could launch thousands of extremely low-cost, battery-operated imaging satellites to obtain a global snapshot of the planet Mars. If the number of elements in the swarm is sufficiently high, the snapshot could be obtained in a matter of hours, provided the throughput of the system is sufficiently fast to transfer the amount of data to a ground station. The total required lifetime of each of these satellites then only amounts to those few hours of operations. Whether anyone would ever be able to justify launching such a large volume of spacecraft to obtain a single image is another matter though, as it will only result in a single snapshot, as well as create a significant amount of space debris.

The aspect of system reliability can also be expanded to other types of distributed systems, yet the most relevant aspect to this thesis is the question of how to design a satellite swarm such that it achieves a given set of requirements. The property of graceful degradation could be exploited in order to reduce the per-element reliability requirements, as hot-spare copies are likely to be present in a satellite swarm. The mission requirements could then perhaps be met with less-reliable swarm spacecraft. Reducing the per-element reliability requirements, whilst an unconventional practice, can reduce the per-element cost. This can be achieved through removing internally redundant systems in each of the spacecraft, as internal redundant systems increase the cost per spacecraft significantly, yet they could prove to be unnecessary in such scenarios.

The overall swarm design is therefore highly linked to the design of the individual spacecraft, as their properties ultimately define the behaviour and properties of the entire swarm. Given the overall swarm requirements then, how is an individual spacecraft to be designed such that the overall swarm will meet its mission-requirements?

This problem can be transformed into the primary research questions for this thesis:

7. *Which definition of a satellite swarm would be the best fit within the category of existing and planned distributed space system architectures?*
8. *Which types of application areas would be best suited for satellite swarms?*
9. *How to design and optimise a satellite swarm such that it achieves a certain mission goal? What effect does graceful degradation have on this process?*
10. *How to design the swarm elements which, when operated as a satellite swarm, ensure the resulting satellite swarm achieves a given mission goal?*
11. *Which element design would suit the OLFAR mission?*
12. *How to design the most basic swarm satellite?*

The methodology used in the thesis varies per sub-topic. Firstly, a literature survey is conducted, after which the various distributed systems used in space are categorised. Satellite swarms are then added to this classification structure, based on their envisaged properties. This leads to a formal definition of a satellite swarm, which is then used throughout this thesis as a reference.

In order to best assess research question 3, an analysis based on Markov modelling is performed to ascertain which element configuration renders a desired operational lifetime. The elements are combined through a k-out-of-m Markov chain, which is then simulated through Monte Carlo analysis for the resulting system lifetime. In addition, this analysis allows characterisation of the system throughput and reliability.

Research question four lead to a design method for individual swarm satellites which has been proposed in (Engelen, et al., 2011). The method relies on iterative design processes and high-level simulations to ascertain the behaviour and effects of a given swarm satellite design on a specific spacecraft swarm. This method is analysed qualitatively, and the differences between the proposed design method and traditional top-down or bottom-up design processes are highlighted.

The OLFAR mission is used as a reference mission, and a first iteration of the element designs is made. The resulting element design is then analysed for its properties when operating in the intended swarm. This swarm is then sized accordingly, aiming for sufficient coverage, as well as sufficient observation time.

1.6 THESIS OUTLINE

This thesis covers a wide research area. In order to focus the research, selected topics are analysed. This section provides an overview of the rationale for their selection, and their location in the document.

Satellite swarms are “new”. In fact, even distributed systems in space in general, other than constellations, are a novelty. This gives rise to quite some confusion between the distinction between the various types of distributed systems, and terminologies are mixed accordingly. In order to allow bounding the research area of this thesis, a workable definition of a satellite swarm is required, as well as its distinctive properties compared to other types of distributed space architectures. This is discussed in Chapter 2, which also discusses the factors enabling the advent of satellite swarms.

Chapter 3 then discusses the design process of a satellite swarm, assuming the properties of the individual elements are known. Satellite swarms can be designed and sized in order to optimise the system's lifetime, availability and overall reliability. These optimisations are treated in this Chapter.

Given that each of the elements in a robotic swarm heavily relies on its interactions with other elements, these should be taken into account in the design processes used to design the individual elements. Adapted and innovative design processes, tailored to satellite swarm elements, are therefore treated in Chapter 4.

In Chapter 5 the OLFAR swarm is applied as a test case for the design processes and methods treated in Chapters 3 and 4 and investigates selected topics, specific to the OLFAR mission. A detailed introduction is given on the items specific to the form of radio astronomy used in OLFAR, after which the OLFAR swarm is sized, and elements of each of the swarm spacecraft are detailed.

Finally, the overall conclusions and findings are summarised and discussed in relation to their novelty in Chapter 6.

2 SATELLITE SWARMS

Monolithic spacecraft can be classified according to their mass or their function. Further distinctions can be made using the target orbit, for example in the case of communication satellites, between LEO- or GEO orbiting satellites, mostly due to the difference in communication latencies and signal directionality. In the case of distributed space systems however, no well-defined distinctions are commonplace, and terms as “constellations”, “formations”, “networks”, “swarms”, and “clusters” are often used arbitrarily, without a clear or even a specific definition. In order to limit the research space of this thesis, a strict definition of the satellite swarms for use in this work is required.

An initial attempt to distinguish between different forms of distributed space systems was made in (Engelen, et al., 2010). The authors defined satellite swarms as: “A spacecraft swarm is a globally controlled cloud of primitive satellites”. This was later refined to Def. 1 (Engelen, et al., 2013):

“A satellite swarm can be defined as a space system consisting of many identical, egalitarian spacecraft, cooperating to achieve a common global goal”.

Def. 1

The key terms in this definition are the fact the satellites are *identical*, and *egalitarian*², and that they *cooperate*. The large number of spacecraft involved will force the spacecraft to remain simple, as complex spacecraft can cause unpredictable (emergent³) behaviour (Engelen, et al., 2011), (Kornienko, et al., 2004), (Rouff, et al., 2004) through their cooperation, which could be undesirable from an operator’s perspective. The large number also allows for the satellites to avoid using complex, fault-tolerant hardware (Engelen, et al., 2012), (Engelen, et al., 2013) as each satellite in fact acts as a hot-spare of any of the other satellites. This leads to the finding that swarm satellites should primarily rely on the redundancy provided by the presence of other satellites in the swarm, rather than on including redundant internal components (Verhoeven & Jongkind, 2003), (Verhoeven, et al., 2011), (Engelen, et al., 2012). Moreover, due to their egalitarian nature, each satellite is by definition fully capable of replacing another satellite.

This notion can also be expanded, in that swarm satellites should rarely include more than a single payload. Limiting the amount of scientific instruments aboard a swarm satellite would reduce the overhead involved with managing multiple instruments, which in turn reduces the development time required to design the control mechanisms, as well as reducing the individual spacecraft complexity.

Given the cost associated with launching spacecraft, the large number of satellites in the swarm will undoubtedly force their design towards smaller, lower mass spacecraft, as otherwise the launch cost of such a system could prove to be prohibitive. Larger, more costly satellites are not prevented from forming a swarm however, yet the economic

² Egalitarian, adj.: *formal*: “aiming for equal wealth, status, etc., for all people”
(Source: Merriam-Webster Learners dictionary)

³ Emergent, adj.: “arising as a natural or logical consequence”
(Source: Merriam-Webster Dictionary)

benefit is questionable. Monolithic, individual satellites can take part in an ecosystem (Verhoeven, et al., 2011) involving a satellite swarm. The monolithic satellite can use its cooperation with the swarm to its advantage, for example by using the swarm as a target identification system, in which the swarm would identify targets requiring closer inspection. That way, the satellite could avoid spending time inspecting uninteresting targets. In this example, everything depends on the ratio of the respective costs: a large satellite's observation time should be more expensive than the cost of launching, maintaining and operating the supporting swarm, as otherwise the ecosystem would not be (economically) viable. The swarm in turn could also service multiple monolithic satellites, spreading the cost.

The notion of low cost should also include operational costs. Managing a single satellite is an expensive undertaking (Saleh, et al., 2004), and managing a whole swarm is potentially even more expensive. Therefore, a high degree of autonomy within the swarm is seen as beneficial (Tripp & Palmer, 2010). In the case of a satellite swarm, micro-management of each individual satellite is deemed futile, as ad-hoc decisions are best left to the satellites themselves (Pincirolì, et al., 2008), (Kornienko, et al., 2004), (Tripp & Palmer, 2010).

The degree of autonomy a swarm design should aim for is such that a global directive should suffice in controlling the swarm in all but the most exotic of scenarios (Bonnet & Tessier, 2007). This will increase the development cost, as well as development time, yet saves on operational cost. This in turn limits the minimum required operational time of the swarm, as a short-lived swarm's development costs should probably remain limited.

The original definition (Def. 1) can therefore be amended into a guideline definition for an ideal satellite swarm:

"The ideal satellite swarm is defined as a distributed space system consisting of many identical, egalitarian spacecraft, autonomously cooperating to achieve a common global goal".

Guideline
Def. 2

Swarms satellites can therefore be characterised through their inherent properties, in that they are *simple, egalitarian, inter-connected, mobile, location aware (though indifferent), autonomous and cooperative*.

2.1 TYPES OF SATELLITE SWARMS

Swarming is a common behaviour of natural, animate systems. Examples are honeybee swarming or foraging, foraging ant colonies and schooling fish. The purpose of this swarming behaviour is different for each type of swarm, as ants and bees use it for foraging (Beekman & Ratnieks, 2000), (Johnson, 1991) or mating (Hölldobler, 1976); whilst fish use it for hunting (Parrish, et al., 2002), for protection (Brock & Riffenburgh, 1960) or gaining benefit from the hydrodynamic properties (Deng & Shao, 2006) of the school or shoal.

By analogy, satellite swarms would also benefit in different ways from the process of swarming. Three cases are reviewed in this section, which are classified into the system of distributed space systems in the next Chapter. The different classes are derived from natural examples, as well as any mission proposals which would apply a similar approach.

The first swarm type treated in this section will be grouped under the name of “satellite clouds”. Satellite clouds can be seen as the primal satellite swarm, as they lack any form of orbit control, yet consist of identical, egalitarian satellites, which can cooperate to achieve a common goal. An example would be the QB50-mission proposal (Gill, et al., 2013).

The other two satellite swarm types treated in this section are foraging satellite swarms, and schooling satellite swarms. These can be considered as more advanced swarms, which subsequently also apply more advanced swarming methods, as well as orbit control. In nature, no other examples of swarming or swarm-like behaviour are known, yet in the case of man-made swarms, perhaps other types and/or behavioural patterns can be imagined. The list of types presented here is therefore not necessarily conclusive.

2.1.1 Satellite clouds

Satellite clouds apply no or very little active orbit control. This implies they will remain in their initial orbit, save for deviations caused by external perturbing forces. In nature, examples can be found in algal blooms (Platt, et al., 2003), (Sverdrup, 1953), for example, or, to an extent, swarms of jellyfish (Zavodnik, 1987). Blooming (Sverdrup, 1953) is a term used in biology to refer to the scenario in which a certain species rapidly reproduces when the conditions are right. An example is an algal bloom, when phytoplankton reproduces where nutrient concentrations are high. They lack significant means of locomotion; hence rapid reproduction is their only means of harvesting most of the nutrient concentration.

In the case of satellite swarms, the phytoplankton could be represented by extremely simplified smart sensor systems. Simplified means of locomotion might allow them to converge near regions of interest, mimicking the behaviour of jellyfish (Purcell, et al., 2000), such as for example differential drag control. Ideally, these sensors would be launched in vast quantities, relying on their large numbers to ensure observations of interest are made within a given time span. Examples of such smart sensor systems would be clouds of femto-satellites, satellites-on-a-chip, or smart-dust clouds (Barnhart, 2008).

Currently the most advanced example would be the QB50-initiative, which aspires to launching a mixture of 50 double- or triple-unit CubeSats into very low earth orbit, performing in-situ sensing of the lower thermosphere (Gill, et al., 2013), as well as the Planet Labs Inc. “flock” of doves, which image Earth at regular intervals (Boshuizen, et al., 2014).

Fig. 2.1 gives a schematic, simplified representation of the acceptance criteria with respect to the relative position of some of the swarm’s elements, using a simple colour scheme. This figure is, due to the indifference or even incapability of satellite clouds to relative position control, not truly relevant. However, since a similar figure will be given for the other types of swarms, it allows direct comparisons between them.




Fig. 2.1: (the lack of) Inter-element position control in a satellite cloud⁴.

Another example of a satellite cloud would be the SURO initiative, in which one of the design options was to deploy a number of radio receivers in deep space, and allow them to drift away (Baan, 2012).

2.1.2 Foraging satellite swarms

Certain swarming insect colonies excel at harvesting resources across vast areas. Examples are honeybees (Beekman & Ratnieks, 2000) and ants (Johnson, 1991), when considering their efficiency at harvesting food supplies from areas surrounding their nests or hives. Their strategy is initially based on random motion, in order to increase the chance of encountering a source of nutrients. Once identified, the workers ensure other workers can find their way to this source, either through the use of pheromones (Evison, et al., 2008), or in case of bees through mapping and communicating the route to take to the food supply (Gould, 1975), until it is depleted. Some random errors are built in to the system, e.g. in the case of worker bees it has been shown that younger worker bees are less effective at decoding the instructions towards a food supply than older, more experienced workers (Seeley, 1983). External factors also introduce certain errors into the acquisition process, ensuring the supply-acquisition process continues, even when a source has been identified. This strategy allows them to efficiently identify and harvest nutrients from every imaginable supply in a given area around the hive.

⁴ The colours represent the acceptance level of the position with respect to the central satellite (black). Red represents an unacceptable position, orange represents a case which could briefly be tolerated, and green is fully acceptable.




Fig. 2.2: Control strategy for a harvesting swarm⁵.

Foraging using purely local communication for robotic swarms has been studied before (Hoff III, et al., 2010), and they find that ant pheromone trails significantly increase the foraging efficiency; and they also propose alternative “virtual” pheromone signalling methods, which increase the efficiency even more. This requires the robots to feature some form of locomotive capability however, which will cost energy and propellant, both of which are expensive.

Foraging satellite swarms are most useful when their distribution is randomised. This is possible, for example through only giving them a minimal distance to the nearest neighbour constraints, rather than both a minimum and a maximum. That way, the swarm elements will naturally try to move as far away from their nearest neighbours as possible, resulting in a mesh-like orbital distribution. The size of the mesh then in turn is determined by the number of satellites launched. Also, given the delta-V requirements involved in changing orbital planes, the mesh will likely be somewhat inhomogeneous, due to the limited inclination changes performed by the satellites. This inhomogeneity can be used as a useable and tuneable parameter however, for example by giving the swarm an incentive to cover select areas more frequently.

The primary application area for foraging swarms would lie in (rapid) global observations, as a maximised distribution in low orbits is most useful in detecting fast-acting phenomena. An example mission is the Fuego fire monitoring constellation (Escorial, et al., 2003), which is shown to be able to benefit from being managed as a

⁵ The colours represent the acceptance level of the position with respect to the central satellite (black). Red represents an unacceptable position, orange represents a case which could briefly be tolerated, and green is fully acceptable.

swarm (Bonnet & Tessier, 2007). Also a global, low latency communication network, not unlike the Iridium constellation would present an application ideally suited for a foraging swarm.

A simplified inter-element position control example is given in Fig. 2.2, showing that each element tries to “push” all other elements away, eventually maintaining a minimal distance between all elements.

2.1.3 Schooling satellite swarms

Schooling is a process which is usually used by animals for protection, although sometimes it is also used for hunting (Partridge, et al., 1983). When a swarm is said to school, they attempt to maintain some form of a geometric distribution; maintaining a distance with respect to the nearest neighbours which is bounded, and generally minimised (Parrish, et al., 2002).

In the case of satellite swarms, schooling swarms can be applied to missions requiring a distinct, bounded, physical distribution of the elements, such as for interferometry missions. A distinction must be made however between formation flights and schooling swarms, in that swarms have no preference for the position of a distinct individual element, as only the global geometry is defined. This means that an optical interferometer requiring a central hub for combining the collected light is *not necessarily* a satellite swarm, as the position of the hub will always be in the central point of the geometry, giving it a specialised function. In case the hub function is performed by a separate satellite, with a different design, this particular satellite would not be part of the satellite swarm. However, the collecting nodes could still be swarm satellites, cooperating with the hub in the form of swarm-hub interaction, much like in a satellite eco-system (Verhoeven, et al., 2011).

The orbital constraints for swarms do not have to be defined as strictly as for example formation flights, as the interchangeability of the elements or nodes allows them to choose their relative orbits such that they can minimise their global or local propellant consumption. Strict position boundaries hamper that process, reducing the effectiveness of the swarm and biasing mission design trade-offs towards a formation flight, yet with proper control system design, it can be achieved at minimal energy expense (Pincirolì, et al., 2008).

As Fig. 2.3 shows, in case of a schooling swarm, the swarm attempts to maintain defined relative distances. In certain cases, even relative angles can be maintained, yet since all elements are functionally identical, the swarm itself is indifferent to which specific individual satellite remains in which specific relative position “slot”.




Fig. 2.3: Control strategy for a schooling swarm⁶.

2.2 TAXONOMY OF DISTRIBUTED SPACE SYSTEMS

While making a distinction between the different types of distributed space systems helps in defining the boundary of the research space, it can also help in identifying distinguishing parameters, as well as suitable application areas for a certain system. Each type of distributed space system will have its merits, as well as unique area in which it excels. This Chapter attempts to identify those, while focussing mainly on satellite swarms.

Four distinguishing properties are considered: the orbital distribution strategy, the constituency, the control methods and hierarchy and the communication methods. A quick glance at potential communication methods however shows these to be mostly irrelevant in distinguishing between the various types of distributed space system architectures.

Mass, which is the distinguishing parameter most commonly used for monolithic satellites, is primarily a cost driver. In distributed space systems, the mass of the individual spacecraft is inferior in terms of cost to the launch costs associated with the aggregated mass of the entire system. Distinguishing systems using spacecraft mass can therefore be deceiving.

⁶ The colours represent the acceptance level of the position with respect to the central satellite (black). Red represents an unacceptable position, orange represents a case which could briefly be tolerated, and green is fully acceptable.

2.2.1 Orbital distribution

Using the modified swarm definition (Guideline Def. 2), combined with the classification attempt made in (Engelen, et al., 2011), a rough taxonomy based on the orbital distributions can be laid out, as is shown in Fig. 2.4, for a homogenous system, i.e. for a system in which all participating satellites are functionally identical (Sabatini & Palmerini, 2009).




Fig. 2.4: Classification tree of distributed space systems (1) according to their orbital distribution strategies, for functionally identical satellites.

In the case of a constellation, such as the various GNSS constellations (Hegarty & Chatre, 2008) and the Iridium constellation (Maine, et al., 1995), satellites are spread into orbital slots, with the number of satellites and orbital planes optimised so as to optimise coverage. Foraging swarms (see section 2.1.2) also adhere to this strategy. The distinction is that in the case of a swarm, the particular satellite covering a certain slot is irrelevant, as all satellites are interchangeable. The notion that swarm satellites could autonomously determine where to go causes the slots to be defined “on the fly” and ensures robustness of the system against the loss of individual satellites, but does not change their objective to maximise their distribution.

Formation flights, as well as schooling swarms (see also section 2.1.3) are instructed to maintain a certain geometric distribution, for at least part of the orbit (Sabot, et al., 2001). In the case of a schooling swarm, interchangeability implies a global propellant optimisation strategy can be used to ensure coverage. Also, the orbital maintenance strategy could be simplified. As an example, two outer limits can be set, ensuring a minimal inter-satellite distance, as well as a maximum. The practical implications are then handled by the swarm, instead of ground-based mission control, or ultimately the end user.

Clouds are a special case, as they apply no orbital maintenance at all. An example would be the QB50 initiative (Gill, et al., 2013), which uses a cloud of 50 nano-satellites using CubeSat platforms to perform in-situ sensing in the thermosphere. Once released from their launcher, they are allowed to drift freely until they re-enter. Another example would be the SURO-LC (Baan, 2012) mission proposal, which used the term “passive formation flying”. SURO-LC does apply a certain degree of orbital control during the deployment phase, but the intention is to avoid orbital control all together.

When considering in-homogenous systems as well, the classification tree can be expanded, as is done in Fig. 2.5. The difference appears to be marginal, as they are all essentially clusters (sometimes also called fleets or convoys) of spacecraft. Clusters differ from eco-systems (Verhoeven, et al., 2011) in that they require cooperation in order to be effective, whilst an eco-system is potentially less reliant on the symbiosis.

Two distinct example cases of heterogeneous systems are the A-Train (Afternoon Train) (Schoeberl, 2002), and fractionated spacecraft (Chu, et al., 2013). The former combines Earth-observing spacecraft, each carrying a different instrument set, while the latter breaks up functions of a single spacecraft into distinct spacecraft. The A-Train can therefore be regarded as a sub-set of fractionated spacecraft, with only a distribution of payloads, and perhaps a lack of a direct inter-satellite communication link. In more distributed fractionated spacecraft also other functions, such as data storage and processing, are distributed and at times, formation control is also distributed (Chu, et al., 2013) through an inter-satellite communication network.

2.2.2 Constituency

Homogeneous distributed systems consist of (functionally) identical satellites, as is the case for satellite swarms (Sabatini & Palmerini, 2009). Distributed systems can also consist of a combination of (functionally) different types of spacecraft, which can then be considered as a heterogeneous distributed space system. Each type of spacecraft in a heterogeneous system can be tailored to a specific task or function, as was done for the A-Train formation flight.

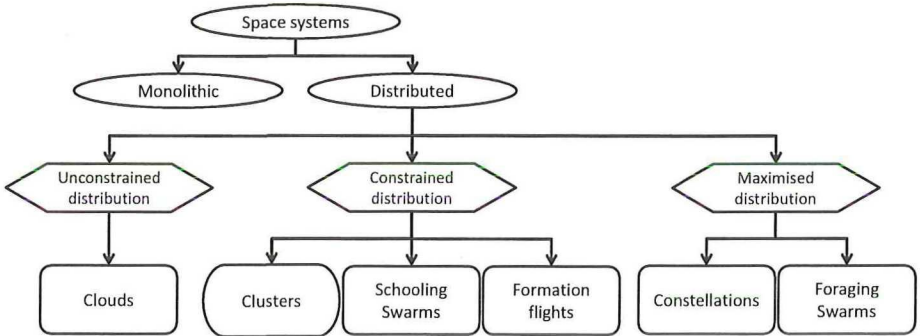


Fig. 2.5: Classification tree of distributed space systems (2) according to their orbital distribution strategies, including inhomogeneous systems.

Clusters and formation flights are generally heterogeneous, although some formation flights can be homogeneous, e.g. SURO-LC (Baan, 2012). Most distributed systems are homogeneous however, as in the case of the Iridium and GNSS constellations. An overview is shown in Fig. 2.6.

2.2.3 Control strategy and hierarchy

Distributed systems can be controlled from the ground segment, with commands sent to each individual spacecraft. This is the most commonly used method, for example for the Iridium and GNSS constellations, as well as for NASA's A-Train formation, as the distribution is important to the operation of those systems, and since these satellites lack

a sufficient degree of on-board autonomy and/or a means of inter-satellite communication to allow coordination of the system in-space.

Research into autonomous control for formation flights, for example for the PROBA-3 mission (d'Amico, et al., 2008), (Sánchez-Maestro, et al., 2013), as well as distributed control strategies (Jafari, et al., 2010), (Massioni, et al., 2008)), (Chu, et al., 2013) for formation control is being performed. In-flight experience in autonomous formation flying strategies has been gained with the PRISMA mission (Gill, et al., 2007), and the TerraSAR-X and TanDEM-X missions (Montenbruck, et al., 2008).

Plans for nano-satellite missions performing (autonomous) formation flying exist; most notably the Can-X 4 & 5 satellites (Armitage, et al., 2013) and TU Delft's DelFFi mission (Gill, et al., 2013).

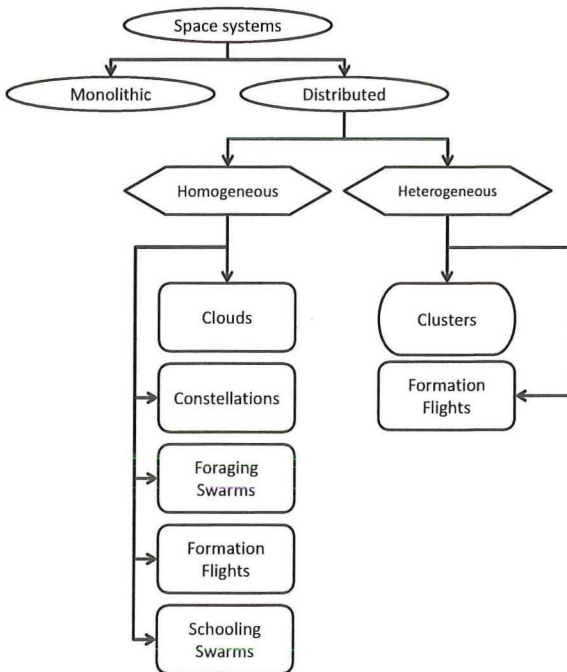


Fig. 2.6: Classification tree of distributed space systems (3) according to their constituency

In the case of formation flights, control is generally global, which implies the algorithms are designed such that the formation requires accurate information on the location and intentions of each of the satellites present in the formation.

For satellite swarms, the control algorithms can be designed as a global control system as well, which manages each of the satellites in the swarm individually. This can either be centralised, in which case one of the swarm satellites commands the others, or preferably distributed (Izzo & Pettazzi, 2007), in which each of the satellites computes (part of) the solution. Research into localised swarming methods (Vos, et al., 2013), (Aso, et al., 2008), (Xue & Zeng, 2008), (Hoff III, et al., 2010), which controls the global swarm

configuration only through local interactions is flourishing. Such mechanisms would eliminate the requirement to transmit the positions and intentions of each of the members of the swarm to all other members, as only nearest-neighbour information is required.

All of the above control strategies are possible for all types of distributed systems; yet only the more modern architectures consider using novel strategies like distributed control and local control.

An overview is shown in Fig. 2.7 (a) and (b). In this figure, the control strategies are linked to their orbital properties, in that a distinction is made between systems limited to their initial orbit, and systems able to significantly change their orbital properties. As can be seen in the figure, satellite clouds are limited in terms of control (they in fact provide no active control), whilst constellations and formations actively control the positions of the individual satellites; either relative to each other, or to an absolute reference frame.

Satellite swarms calculate the geometric or gravimetric centre of the swarm, and the satellites position themselves relative to it. This then allows operators to control the properties of the geometric centre, rather than each individual spacecraft.

Another distinctive property lies in the hierarchy of the system. In formation flights and satellite clusters, the order of the satellites is important, as the satellites are not functionally identical to one another. Traditionally, satellites in constellations are given unique identification numbers to allow ground-based control, which limits their flexibility, and hence renders a hierarchical system. Certain formation flights define a "master satellite", which controls the formation, as is done for example in the DARIS study (Saks, et al., 2010), the SURO-LC proposal (Baan, 2012), as well as ESA's Darwin mission (Rabbia, 2004), or NASA's Terrestrial Planet Finder (TPF) (Beichman, et al., 1999). These satellites perform a special function, and are therefore unique, which renders them irreplaceable.

For satellite swarms, following the definition, no central satellite can (strictly) be defined. Even if a "master" satellite is required, for example for centralised payload data correlation, satellites can take turns in performing the role of master. This makes a swarm an egalitarian system, which in turn eliminates strict requirements on orbital control. For a satellite swarm, as long as data at a certain location is gathered within a select time-frame, the actual satellite performing the acquisition of the data remains anonymous and irrelevant to the cause.




Fig. 2.7: Control strategies of distributed space systems (a)

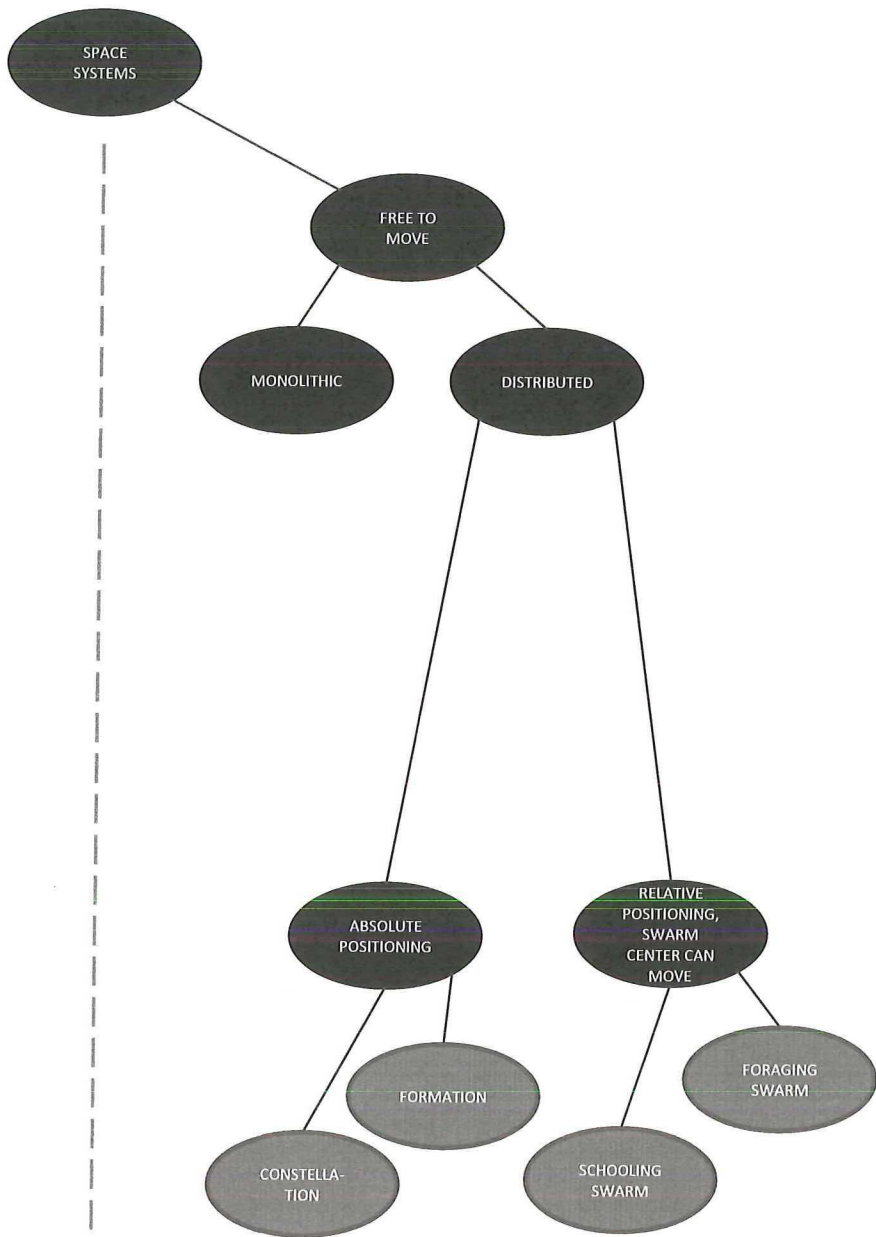


Fig. 2.7: Control strategies of distributed space systems
(b)


2.2.4 Communication methods

Distributed systems require communication, in order to allow for data collection and data aggregation, as well as allowing for command and control operations. The methods of performing communication can differ quite significantly between various types of distributed systems, even though certain similarities can be used in order to group them more effectively. An overview of the different methods which were distinguished is shown in Fig. 2.8.


Today, the primary means of communication with spacecraft is through a direct link with a ground station, or a network of ground stations. This method can be used in a distributed system as well, if these links can also be used to relay information between two or more spacecraft. This is shown in Fig. 2.8 B, and termed “ground-station-in-the-loop communication”. Satellite clusters for example rarely feature a direct inter-satellite link. This requires them to each transmit their data to their ground station, which can in turn correlate the data for further processing. Data from the other satellites in the cluster is gathered independently. This can cause a significant lag in the availability of the data, which is why the Iridium constellation for example featured an inter-satellite link.

Spacecraft within the distributed system can also use an inter-satellite link (when present) to relay information directly. This is called point-to-point communication, as is shown in Fig. 2.8 A. In case of point-to-point communication, no networking protocols are in place, and therefore, all spacecraft will have to establish a direct communication link with each of the spacecraft they have to communicate with. This is more often than not the case for formation flying satellites, which can then use their inter-satellite link for performing in-space data correlation, as well as for low-latency communications, for example for fine-grained control of the formation. Some formations (Kahle, et al., 2012), use a ground-station-in-the-loop scheme to either control the formation, or to transfer data between the satellites’ orbit control computers.


Due to the large number of spacecraft involved in satellite swarms, they have to coordinate their actions almost continuously, which rules out the use of a ground-station in the loop in almost all but the most exotic of swarm missions. The communication method however can differ significantly. Swarms are a paragon of the potential of peer-to-peer communication protocols and nearest neighbour-only networking, shown in Fig. 2.8 D. Global (swarm-level) distribution of data is, in the case of nearest neighbour-only communication, possible through for example gossip (Leitão, et al., 2007) or flooding protocols (Paruchuri, et al., 2003). Swarm satellites performing peer-to-peer networking (Qiu & Srikant, 2004), in which case each satellite maintains an active copy of the data present in all other satellites is also possible, but quite likely much more demanding.




A: Point-to-point communication



B: Ground station in-the-loop communication



C: Networked communication



D: Nearest-neighbour communication

Fig. 2.8: Communication principles in distributed systems

Rerouting of signals inside the swarm, as is done in terrestrial TCP/IP networks (Autenrieth & Kirstädter, 2000) is also possible (Budianu, et al., 2011). In this case, each satellite has a unique network address, and data destined to a specific satellite is transmitted through the network through a scheduled route, as is shown in Fig. 2.8 C. Networked communication can give satellite swarms, and indeed also satellite clouds an extremely high degree of flexibility in terms of choice of communication strategy, as the routing can be determined “on-the-fly”.

An overview of the distributed architectures and their methods of communication is shown in Fig. 2.9. Noteworthy is that not one particular communication strategy is solely used by a particular distributed architecture. However, trends are visible, in that highly distributed systems are more likely to prefer networked or nearest-neighbour schemes, in order to limit the latencies, as well as the load on the ground station.

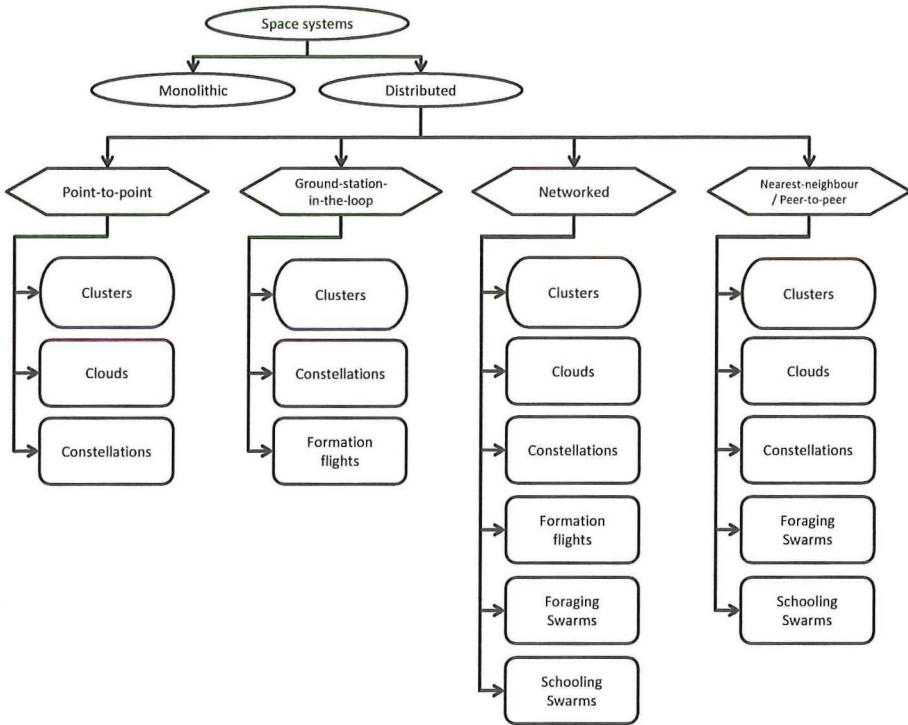


Fig. 2.9: Classification tree of distributed space systems according to their communication strategies

2.2.5 Synthesis

All of the previously identified traits can be combined into the overview graph shown in Fig. 2.10. Communication schemes are left out, as they do not appear to be a distinguishing property.

One apparent void occurs in constrained systems, with a one-side limited orbital distribution, and a heterogeneous constitution. These would show behaviour similar to Foraging swarms. Only one reference to such a system was found, in an Earth-observing eco-system (Verhoeven, et al., 2011), although more situations could be imaginable.




Fig. 2.10: Classification tree of distributed systems according to their orbital distribution, their constituency and their control strategy

2.3 SWARM MANAGEMENT

Managing a cloud is a rather uncomplicated process, as cloud satellites do not necessarily interact with each other. When adding a network layer between the nodes to allow routing data from elements to a single ground station, management complexity increases a little, yet the return is that ground station operations are then somewhat relieved, as passes involving multiple satellites can for a large part be ignored, as the data can be routed through each of the contacted spacecraft.

Schooling swarms are much more complex systems however, with many agents orbiting in close proximity. Planning, synchronisation and control of the observations can be an elaborate process, requiring a lot of communication overhead. Foraging swarms are equally complex, in that careful planning of their individual observations is required, in order to guarantee a certain degree of coverage over select areas, and due to the intersections in their orbits, high velocity collisions have to be avoided through careful orbit planning. The communication overhead involved in such planning efforts can be significant, and moreover, the distance between each of the satellites in a foraging swarm is significant by design, so communication delays are inevitable. In case a swarm is over-defined, task allocation can be scheduled such that the global energy levels are optimised, as was studied by Liu et al. (Liu, et al., 2007).

Directly managing a swarm from a ground station will therefore be a complex process, which is why it is often suggested to control swarms through stimergetic means (Tripp & Palmer, 2010), or through other forms of global incentives (Pincioli, et al., 2008), rather than direct, individual commands. That way, satellites receiving a directive can inform the rest of the swarm; ensuring all elements are informed of the change. While this is quite promising in significantly reducing the involved managerial load and overhead (Bonnet & Tessier, 2007), (Izzo & Pettazzi, 2007), it is also a novel area, and quite some research is still required in order to guarantee the robustness of such control methods.

While the configuration and hence orbital distribution of swarms can be controlled through decentralised means and incentives, the global orbital locations of for example schooling swarms can be controlled through passing the location of the geometric or gravimetric centroid of the configuration, as all individual satellites are interchangeable. The swarm should then be able to optimise the propellant consumed in order to achieving the requested configuration.

Crucially, data gathered by satellite swarms can be pooled. If a swarm acts as an imaging cloud, the communication bandwidth required to transmit each individual image down to Earth will be significant. It can then make sense to temporarily buffer the data in the swarm, and provide it on demand, as it is quite likely that most of the captured imagery would be thrown away in case there's no specific interest in that region. More elaborate schemes can also be envisaged, involving data compression or schemes in which the swarm satellites store a reference image, and only report observed differences. Swarms can also be given more autonomy in determining which images are important enough to send down; for example in the case of forest fire detection, in which case a certain temperature threshold could be set which has to be met in order for the swarm to report the occurrence.

In case of certain schooling swarms, for example the OLFAR mission (Rajan, et al., 2011), data from each individual satellite is correlated in space, after which the result of the correlation is sent down to Earth for further processing. This leads to a reduction in the data volume, as only a single, correlated data stream is sent down to Earth, instead of

'n' individual streams, with 'n' representing the number of active satellites. The users of this particular swarm however will have to accept that in doing so most of the data is irrevocably lost. This concept will take some getting used to by the scientific community, and it will occur quite frequently for other swarm missions as well, as the data rates generated by satellite swarms will easily congest all of the available bandwidth. This implies that either a pre-selection of the data to be transmitted (e.g. an "on-demand" service), or significant pre-processing or compression in space will have to be used to reduce the bandwidth requirements. The trade-off therefore will quite likely lie in data volume versus data quality. This does not imply that the data from a satellite swarm is of lower quality by definition, yet it will be more difficult, if not impossible, to verify the data quality after the initial or intermediate calibration periods, if any pre-processing was applied to the data.

2.4 ENABLING FACTORS

The concept of satellite swarms is new. Even highly distributed space systems only appeared in papers starting in 1984 (Molette, et al., 1984), after which it still took many years for technology and other factors to warm up to the concept. Spacecraft swarms are even more of a novelty, with the first references dating back to the year 2000 (Curtis, et al., 2000), as it took the convergence of certain enabling factors, which all played a role in allowing the concept of spacecraft swarms to emerge. Those are the advent of potent and low cost mainstream technologies, a standardised (and widely accepted) platform, and advances in computer sciences, which allow for a new architecture, as well as a paradigm change.

2.4.1 Technology: Spin-in

Mainstream technologies surpassed the capabilities of space-grade technologies a long time ago. Modern mobile phones have access to more storage space than any satellite has seen to date, and they have much more processing power than any space-grade processor on the market can offer. This situation will persist, as the cost of a modern IC factory (called "fab") is well over one billion euro (Ginosar, 2012), which is more than the total space budget of many nations. It is therefore highly unlikely that any space organization will construct or modify a fab specifically tuned to space-grade technologies, especially given that the number of units produced would never reach the numbers produced in mainstream industries (Ginosar, 2012). Also, special space-grade processes, such as silicon-on-sapphire, have been discontinued in most fabs, as well as almost all production facilities for process nodes larger than around 350 nm, which implies that established space-grade IC designs cannot be produced anymore.

The effects are twofold:

- 1) Space grade components which are fully radiation hardened have become extremely expensive to (re-) produce, and are therefore unlikely to be chosen for low cost missions. Space missions are now using up the remaining stock of previously produced components. This stock is dwindling however, and this also hampers all forms of innovation.
- 2) Commercial (non-space grade) components are being used or earmarked for use in low cost missions. Their performance is much higher than their space-

grade counterparts, if they even exist, yet their radiation tolerance is questionable, as are their packaging and solderability. Also their production periods are short, as the commercial market moves much quicker than the space market, causing components to become end-of-life either during the mission, or even during development. This implies that spacecraft cannot use direct heritage of the previous generation, as those components have become extinct. This will cause frequent re-designs of the hardware, which in turn can lead to improvements, but will increase the amount of testing required in order to validate the new components. Alternatively, due to the relatively low cost of the individual components involved, large quantities of subsystems can be produced and placed on the shelf, for later use. When stored properly, this could result in a stock useful for at least a few generations of spacecraft.

One important note should be that even though the cost of commercial components rarely exceeds a few tens of euros per component, it can still take a significant effort to design a circuit with them which will withstand the space environment. Also, commercial components are tested to different standards and for different applications, which generally forces re-certification of these components.

In case of spacecraft swarms however, mainstream technology offers an advantage in that their qualification procedures and more importantly, their high production volumes guarantee consistency: each produced device will show a performance close to the next device, and individual performances are guaranteed by the manufacturer, hence all values quoted in the datasheet of a device will at least be met. In case of the production of a satellite swarm, with almost identical satellites, it would therefore be possible to certify and qualify only a single spacecraft. All other spacecraft in the same series or batch can then be subjected to simple functional and acceptance tests, instead of going through a full test cycle for each of the spacecraft, which in turn reduces the overall cost.

2.4.2 Platform: Standardisation

Nano-satellite numbers have seen a boost compared to all other satellite mass categories. When categorising small satellites according to their mass and launch dates, as shown in Fig. 2.11, a steady increase in numbers of satellites launched can be seen starting from 1957, with most of them being micro-satellites. Note that the former Soviet Union communication satellites of the Strela-1 and Strela-1M satellites are shown separately, in order to highlight individual satellite platforms. When miniaturisation efforts finally came to fruition around 1988, small satellites became more accessible, and more importantly, due to the advent of readily available micro-electronics, more potent, causing a rise in their acceptance. Nano-satellites numbers have only recently started booming, but their numbers are still increasing (Swartwout M. , 2012). Even when factoring in their reduced launch costs, their numbers, and more importantly, the rise in their numbers still exceeds that of the small satellite platforms of the 1990's. This can in part be attributed to the increasing availability of even smaller micro-electronic circuits with ever increasing degrees of integration, and perhaps more importantly, ever increasing flexibility. Yet when looking at the type of nano- and pico-satellites launched since 2003, almost all pico-satellites (save for the two subsatellites launched by the Japanese IKAROS solar sailing demonstrator) were based on the CubeSat platform (Janson, 2011).

The CubeSat platform, introduced by Robert Twiggs (Stanford University) and Jordi Puig-Suari (California Polytechnic State University (CalPoly)) in 1999 (Shiroma, et al., 2011), is a standardized specification for pico-satellites, which includes volume, mass, materials and operational restrictions, and has since grown into a widely accepted platform. The basic building block is a 10x10x10 cuboid structure, which is called a “unit”, or “U”. Multiple units can be combined to form larger satellites. Given that the mass is essentially constrained, initially to 1 kg per unit, and later on to 1.33 kg per unit (California Polytechnic State University, 2009), this implies that aggregate CubeSats, employing multiple units, fall into the category of nano-satellites by definition, or even into the category of micro-satellites. A few examples are shown in Fig. 2.12.




Fig. 2.11: Number and mass of small satellites launched since 1955 adapted from (Janson, 2011)




Fig. 2.12: Various CubeSat compatible satellites⁷.

In that figure, (a) is a 1U CubeSat called SwissCube (Noca, et al., 2009), which is notorious for its attitude control system malfunctioning due to an under-sampling condition. This was solved (Overlack, et al., 2011) however, and the satellite is now functional. (b) is a depiction of UKube-1 (Harriss, et al., 2011), launched in July 2014, which is a mission to test new technologies, including attitude determination and control, as well as a GPS device which will measure plasmaspheric space weather. (c) is a render of NASA's O/OREOS (Nicholson, et al., 2011) satellite, which is an acronym for Organism/Organic Exposure to Orbital Stresses, is a satellite examining the effects of the space environment on organic molecules and biology. (d) is an image of the CSTB-1 satellite (Taraba, et al., 2009), built by the Boeing Corporation, which is another technology demonstrator mission, aimed at, amongst others, maturing commercial low power processors, CMOS ultra low power imagers and associated software algorithms. (e) is an image of the Radio Aurora Explorer 1 (RAX-1) satellite (Cutler, et al., 2010), launched jointly with O/OREOS and other CubeSats in 2010. It operated for only two months, after which the solar panels degraded prematurely resulting in a loss of power. RAX-2 has been launched in 2011, and has been operational since. The mission is aimed at examining the physics behind the formation of magnetic-field aligned plasma irregularities (FAI), which are known to disrupt communications with spacecraft. (f) is an image of the Delfi-C3 satellite (Ubbels, et al., 2005), built at the TU Delft, and launched in 2008. It is a technology demonstration mission, examining the behaviour of thin film solar cells in space, as well as an autonomous wireless sun sensor. The satellite has been operational ever since.

Internally, many CubeSats have adopted an adapted version of the PC/104 PCB standard (PC/104 Consortium, 2013). Out of all satellites shown in Fig. 2.12 however, only the RAX-1 and UKube-1 satellites have followed this standard for their payloads and primary systems. This is mainly due to the lack of commercially available

⁷ Image credits: (1) EPFL, (2) Clyde Space, (3) NASA, (4) Boeing Corporation, (5) The Michigan Exploration Laboratory, (6) TU Delft

components at the time of development of most of these satellites. Recently however, many space companies and also quite a few start-ups have started offering off-the-shelf components (see for a few examples Fig. 2.13). Due to the high degree of standardization, almost all of those components are compatible, which in theory would allow for a very rapid development. The Air Force Research Laboratory (AFRL) is developing a system for adapting “plug-and-play” approaches for use in space. Theoretically, this would allow reducing the development time of a satellite from months to days. Such a feat would not be possible without a significant degree of standardisation, as well as applying some form of plug-and-play compatible operating system running on the satellites, as well as on ground (Lyke, et al., 2005).

Standardisation and low masses also have an effect on the overall cost. Since components are available off the shelf, development costs are in fact shared across various missions. It also allows for “mass-production” of the individual components. Combined with the low mass of nano-satellites, and the reduced qualification procedures involved due to the use of a standard deployer (Chin, et al., 2008), the overall cost of a nano-satellite mission is significantly reduced.




Fig. 2.13: Various CubeSat compatible components⁸.

⁸ Image credits: Clyde Space, www.cubesatshop.com

2.4.3 Architecture: Increased autonomy

Since the advent of modern microprocessors, starting in 1971 with the launch of the Intel 4004, commercial processor performance has been increasing steadily. Gordon E. Moore projected (in 1965) (Moore, 1998) that processors would double in number of active transistors every two years, which would also lead to a performance increase. This projection has been used as a guideline for the manufacturing industry, which resulted in his projection remaining true throughout most of the years following his prediction.

An overview of a selection of processors used in space projects is shown in TABLE 2-II. As can be seen from the table, processors used in spacecraft however have not followed this trend as closely as their commercial counterparts. In fact, micro-satellites were primarily called micro-satellites due to the increased usage of microprocessors in their critical subsystems (Sweeting, 1992), which allowed them to perform much more complex tasks compared to their more traditional (programmable) sequencer-based counterparts. Nano-satellites, due to the reduced financial risks involved⁹ are using more modern processors than any traditional spacecraft would ever attempt, as is shown in TABLE 2-I.

TABLE 2-I
OVERVIEW OF PROCESSORS USED IN CUBESAT MISSIONS

<i>Processor</i>	<i>Performance</i> ¹⁰	<i>Cubesat Mission</i> ¹¹
TI MSP430	2.3 DMIPS 8 MIPS	Delfi-C3, SwissCube, Hawksat, Delfi-n3Xt
Microchip PIC16F877A	5 MIPS	Hayato
Atmel AVR	5.3 DMIPS 16 MIPS	AubieSat-1
Renesas H8/300, H8S	18.9 MIPS	Cute-1, UWE-2
ARM7TDMI	60 DMIPS	BEESAT-1, SwissCube, Jugnu, PW-Sat
Atmel UC3A0512	91 DMIPS	StudSat
Marvell PXA270	780 DMIPS	RAX-1, RAX-2
TI OMAP 4460	6000 DMIPS	STRAND-1

⁹ With risk defined as Risk = likelihood of occurrence * impact

¹⁰ MIPS (Million Instructions Per Second) are used primarily for microcontrollers, or processors for which no Dhrystone benchmark results were found.

¹¹ Source: (Klofas & Leveque, 2013)

TABLE 2-II
TIMELINE OF PROCESSORS USED IN SPACE PROJECTS AND IN COMMERCIAL APPLICATIONS

<i>Year of introduction</i>	<i>Space-grade processor</i>	<i>Performance</i>	<i>Commercial processor</i>	<i>Performance</i>	<i>Example missions/ applications</i>
1974	NSSC-1	Unknown			Landsat-D
1974			Intel 8080	0.330 MIPS	Space Invaders Arcade Games
1976	RCA1802	0.1-1.2 MIPS			Space Shuttle, Galileo
1975			MOS 6502	0.5 MIPS	Commodore 64, Apple II, Nintendo NES
1980	1750A	0.5-3 MIPS			Cluster, Rosetta, Cassini, Clementine, Envisat
1982			Intel 80286	2.66 MIPS	IBM PC
1995	RAD6000	35 MIPS			Spirit and Opportunity, Deep Space 1, MESSENGER, STEREO
1996			Atmel AVR	16 MIPS	Arduino
1996			Intel Pentium Pro	541 MIPS	IBM PC
1997	Mongoose V	16 DMIPS			EO-1, MAP, CONTOUR, New Horizons
1999			Intel Pentium III	2054 MIPS	IBM PC
2000			AMD Athlon	3561 MIPS	IBM PC
2001	RAD750	266 MIPS			Deep Impact, MRO, Kepler, WISE, Juno
2002			ARM 11	1000 MIPS	Samsung S8000 Jet mobile phone
2011			Samsung Exynos 5250	14000 MIPS	Samsung Galaxy S4 mobile phone
2011			Intel Core i7 EE 3960X	177730 MIPS	IBM PC



Using commercial components increases the chance of the occurrence of single-event upset related issues, yet the increased processor performance also implies the time lost due to the downtime following an upset is also minimised, due to the reduction in boot times.

The mobile phone industry in particular (and by extension also the tablet market) has resulted in very powerful embedded processors, with significantly reduced power consumption (Miller, 2012). Perhaps more importantly, these processors are produced in large quantities, which results in very low per-unit prices, and a large user-base of developers and widespread software support packages. The large user-base ensures knowledge about errors and bugs in the devices is widely available, and patches and solutions to the bugs, if available, are generally well-known (Chou, et al., 2001). Mobile operating systems can therefore be considered as robust, in part due to their open-source nature. The use of flash-based memory also allows for in-place firmware upgrades, which, due to the application of a boot loader and dual-memory stores can be designed in a safe way, also tolerant to malicious attempts (Nilsson & Larson, 2008). Reliability figures of mobile phones are not that widely available, although some comparisons can be made. See for example (Jary, 2013). Given the numbers of units involved, it can be assumed that the failure rate of mobile phones is low, as otherwise the repair costs, as well as the damage to the image of the manufacturer would be excessive. Given the complexity of the circuitry inside, and the abuse often experienced by the device, the per-component reliability is therefore likely to be substantial.

It is these properties of high performance, high efficiency, low unit and development cost, wide support system and high reliability which are particularly of interest to nano-satellite builders, as they are able to accept the increased risk of applying unproven technology.

With an increase in available processing power and memory bandwidth and storage, an increased firmware complexity can be handled by the devices. On Earth, such processors are being applied on platforms for research into fully autonomous robotic swarms (Rubenstein, et al., 2012). These robotic swarms are shown to be capable of handling complex tasks, even with limited local intelligence (Zhang, et al., 2007), (Rutishauser, et al., 2009), and are envisaged to be used in dangerous situations for cleaning up mine-fields or creating a robust communication-network on the battlefield, where the low unit cost is beneficial (Şahin, 2005) or in search-and-rescue operations, where a fast and wide-area coverage can be crucial .

In satellite swarms, an increase in autonomy could potentially save a significant part of the overall mission cost, due to the reduction in operational costs incurred by ground station operations. However, the swarm elements will require a robust set of behavioural guidelines, in order to prevent issues with lock-ups. In return, the increased autonomy reduces decision lags, allowing for more accurate ad-hoc decisions, or simply more informed decisions due to the increased data volume transmitted between the nodes making the decision.

2.4.4 Paradigm shift

It will take time before swarm satellites will be accepted in the space community. They show promise for certain applications, mainly due to their reduced per-unit cost, increased reliability and intrinsic large-area coverage.

There are also downsides, as not every type of architecture is suited to a certain task. Satellite swarms will therefore, even after their acceptance, occupy a niche in which a satellite swarm outperforms other types of architectures.

The primary downsides are:

- Space debris: Launching a large number of satellites, combined with a possibly reduced per-satellite reliability figure, will result in a lot of space debris. Not only the swarm satellites themselves, when defunct, will have to be taken care of, but also the deployment systems and upper stages of the launchers will have to be disposed of. Possible solutions lie in launching satellite swarms into very low orbits for example, which would exhibit some form of self-cleaning.
- The risk of unpredictability: Complex, cooperative systems show emergent behaviour (Rouff, et al., 2004). This emergent behaviour can make an anthill into a highly efficient society (Johnson, 1991), but emergent behaviour can also cause disruptions in power-lines and networked systems (Rinaldi, et al., 2001). In case of an autonomous satellite swarm, it has the potential to disrupt operations, or even cause a communications deadlock, possibly resulting in the loss of the entire swarm. Natural swarms have found ways around such lock-ups however, through implementing balance-based decision making processes, which honeybees use for example when choosing a new nest site (Seeley, et al., 2012), (Niven, 2012). Emergent behaviour is currently being studied, and methods to prevent lock-ups have been identified (Kornienko, et al., 2004), yet any newly designed satellite swarm will have to accept qualification tests testing for known issues, gathered from simulations and real-life applications.

The biggest hurdle to take in order for swarms to gain acceptance however is the difference in philosophy. In satellite swarms, unit losses are considered acceptable, which goes directly against the current mentality, which builds upon years of satellite builders and mission designers trying to prevent losses altogether.

The notion in satellite swarms is that the unit costs can be reduced, at the expense of accepting a reduction in the *per-unit* reliability. The *overall* swarm however should provide, through the large number of satellites involved, sufficient redundancy for the system to gain an acceptable degree of reliability and also availability. Availability remains crucial however, as a swarm which remains active, yet fails to retrieve data at the instant it is required does not have any merits.

Increased autonomy, to the point that satellites make their own decisions on nearly every aspect of their mission, is also radically different from most, if not all, missions flown or even considered to date. The processing power available to satellite swarms (using nano-satellite-derived technologies) would allow for full, system-wide autonomy, in which each of the nodes makes decisions on its own orbit control and science observation planning, all in convergence with the plans of the other satellites in the swarm.

In order for satellite swarms to ever come to fruition therefore, a pioneering mission will have to be built and launched. This mission then will hopefully serve as the seed which will enable the paradigm shift required for satellite swarms to gain acceptance as a viable solution to problems difficult to solve otherwise.

2.5 SATELLITE SWARM APPLICATIONS

Satellite swarms feature high numbers, and a high degree of autonomy as a result. The units should remain low cost, in order for the swarm to remain economically feasible. Therefore, applications for satellite swarms will have to exploit one of these properties, as otherwise other architectures will undoubtedly be better suited to that particular task.

Two swarm missions are known to be under development at the time of writing, which are the OLFAR radio telescope (Rajan, et al., 2011), and the ISIS AIS swarm (Verhoeven, et al., 2011). Even though for these particular missions a swarm configuration is deemed the most favourable, *no true killer application for satellite swarms has been identified*. For the OLFAR radio telescope for example, previous studies have shown that either a formation of traditional satellites can achieve similar results (Saks, et al., 2010), (Baan, 2012), and either a single antenna or a traditional radio-telescope positioned on the back-side of the moon (Zarka, et al., 2012), (Klein-Wolt, et al., 2012) could even surpass the observational quality offered by a space-borne distributed system. The advantage a nano-satellite swarm would offer in case of OLFAR however is a reduced unit cost, which in turn allows for an increase in the number of antennas. This then reduces the observation time required, or alternatively increases the sensitivity of the instrument. OLFAR also has access to much higher performance components, compared to DARIS for example, which relies on proven, space-grade technologies (Baan, 2012). The output of OLFAR is therefore expected to be significantly increased compared to DARIS.

The ISIS AIS swarm on the other hand would be equally feasible using a constellation, or even a satellite cloud. Features like orbit maintenance and control aren't required for this particular application. The only significant advantage of a swarm would be offered through the inter-satellite link, in which case the satellite-agnostic property of a satellite swarm would allow any spacecraft passing over a ground station to transmit data gathered by nearby satellites.

The NASA ANTS concept design focuses on asteroid-belt investigations (Curtis, et al., 2000), applying a large number of satellites to increase the chance of an encounter, as well as the effectiveness of the scientific operations during such an encounter. Atmospheric science requiring global, rapid surveillance would also be a possibility, and the QB50 proposal in fact applies a satellite cloud to study the upper atmosphere (Gill, et al., 2013), which indeed seems to be a very viable architecture for this particular mission. Adding orbit-maintenance to each of the spacecraft would allow longer-duration observations, yet also increase the unit-cost.

A few general trends can already be distilled however:

- Swarms have the potential to allow launching into very low Earth orbits, which would allow rapid revisit times, as well as guaranteed disposal after the mission or in case of the loss of an element, due to the high drag forces experienced by such satellites. Satellites in the swarm with defective orbit control capabilities would therefore naturally deorbit in a very limited time-span. It therefore seems likely, and also recommendable to launch satellite swarms into unstable, and also normally unfavourable orbits (Verhoeven, et al., 2011), in order to reduce the space debris quantity. Satellite swarms have the potential to excel in these locations due to the inherent robustness in the system, as well as through their low cost units, which reduce the impact of the loss of an individual swarm

satellite (Tripp & Palmer, 2010). Incidentally, lower orbits allow more detailed inspection of ground-based targets, due to the reduced range.

- Highly distributed missions, especially with very large numbers of satellites and relaxed position-control requirements would benefit from the autonomous control generally considered for satellite swarms, due to the reduction in ground station operations (Tripp & Palmer, 2010). This also holds for the OLFAR swarm for example, in which case also the egalitarianism of swarms allows for continuous upgrades to the number of satellites in the swarm, improving the sensitivity of the system with each upgrade.
- Missions searching for highly dispersed or fast phenomena would also benefit from using satellite swarm architectures. The NASA ANTS concept proposed swarm satellites to find smaller asteroids in the asteroid belt, and also missions searching for forest-fires or earthquakes might benefit. Search-and rescue operations can use a swarm's inter-satellite link to relay information to ground stations. The reliability of the swarm spacecraft however can come into play however.

Mission designers considering a satellite swarm for their application should realise which of the unique properties of a satellite swarm would be beneficial for their particular mission, after which these should be exploited.

3 DESIGN OF A SATELLITE SWARM

Many non-ideal satellites cooperating in an actual environment, with physical limitations on the amount and latency of communications, as well as any limits imposed by the accuracies or timeliness of their sensors causes a phenomenon commonly referred to as "emergent behaviour" (Rouff, et al., 2004) to occur. Emergent behaviour can be defined as behaviour exhibited by a group of cooperating individual (swarm) elements which was not specifically designed for, and which only emerges when elements interact. Emergent behaviour can have beneficial properties, yet due to its unpredictable nature, such behaviour can be unwanted. Also, not all emergent behaviour is benign, as it can for example cause the satellite networks to become unresponsive in reaction to an element continuously transmitting faulty messages. Creating a dependable, reliable swarm is therefore not a straightforward process (Winfield, et al., 2005). Progress has been made towards identifying all possible emergent behaviour for a given swarm, which could be used to verify the emergent behaviour will not cause mishaps for that given swarm (Winfield, et al., 2005). Proving that the identified behaviour is indeed all emergent behaviour which will occur however requires extremely well defined mathematical models of the actual satellites, as well as the environmental interactions.

A swarm relies on cooperation. Its design is therefore also highly dependent on it. However, the inter-element communication can be intermittent due to either availability or even reliability issues with the satellites or loss of line-of-sight due to the distances involved or blockage by other objects. Communication over large distances also involves communication delays and lags, complicating direct control.

Due to these factors, design of a swarm cannot be performed through a traditional straightforward top-down process, in which the desired element behaviour is defined in advance, and the elements are designed to match, as inter-element interactions will cause emergent behaviour which may prove to be harmful to the overall requirements specified in down by the top-down process. Simulations, preferably with hardware in the loop are required in order to tune the design such that the resulting swarm behaviour is considered manageable. In certain satellite swarms, data from the individual elements can also be combined to allow extraction of more or higher level data through for example correlation of the individual datasets. This implies that the element design in turn can in select cases be simplified, as the data products generated by the swarm will be more than the sum of the output of the individual element. Phenomena which are unobservable based on the performance of the individual instruments present in each of the swarm satellites could become observable through interaction within the swarm, which in turn could allow satellite swarms to use lower-cost and perhaps smaller sensors.

This is the precept of the OLFAR swarm, which intends to observe faint radio signals with very long wavelengths, in the order of 10 to 1000 m, with an instrument sensitivity of 65 milli-Jansky (Engelen, et al., 2010). A directional antenna with sufficient resolution at these wavelengths would require an antenna with an equivalent diameter of 100 km which would be impossible to construct with current-day technologies and budgets. Using a swarm of satellites, each equipped with omnidirectional antennas, a virtual telescope can be formed which sports a diameter of 100 km if a number of satellites in the swarm is spaced at 100 km apart. The required sensitivity can be reached through a process of integration over time.

Swarm satellites are not necessarily more reliable compared to traditional satellites however as designers will likely be forced to abandon space qualified (and therefore reliable) hardware for modified COTS hardware in order to decrease per-element, and with it the overall system costs. This can have effects on the availability of the individual satellites, yet it does not have to affect the overall reliability of the system. (Engelen, et al., 2012), (Engelen, et al., 2013)

There are many reasons why the adaptation of commercial electronics in spacecraft is considered a risk, such as quality assurance, short product lifecycles and potential production variances. One immediately apparent issue with using commercial electronics in space is their susceptibility to Single Event Upsets (SEU). In the event of a charged particle entering a silicon device, it can deposit a charge at an arbitrary location. This location can reside in the active part of the device, in which case, the charge can prove to be sufficient to change the state of that part of the device. In digital circuits, memory and registers of ASIC's and microprocessors are the most susceptible. In case of an FPGA, which most closely resembles a memory device, the entire functionality of the effected part of the chip can be altered. Traditional space hardware has generally been "radiation hardened", which is a process involving changing the design of the devices such that they become practically impervious to the effects of radiation and charged particles. Commercial electronics does not apply such practices, hence for those devices a strategy of radiation-tolerance has to be followed. Radiation tolerance entails designing the circuit or the software such that single event upsets are ignored, or otherwise addressed. In case of FPGA's, a process of "scrubbing" is used, which replaces the current code and circuitry with a "golden" copy, stored in a radiation hard device, at regular intervals (Carmichael & Brinkley, 2006). ASIC's and microprocessors can apply Error Control Codes (ECC) to verify the information they receive is valid. ECC checking has been proven to significantly decrease the susceptibility of a given device to SEU's (Finn, 1989), and is currently implemented even in common mainstream devices for various other reasons, such as increasing the reliability of high speed high density memory devices (Normann, 1998), and preventing instability caused by memory leaks (Qin, et al., 2005). Commercial electronics are also susceptible to other radiation and charged particle effects, such as latch-up and single-event functional interrupts, each of which requires different design practices to ensure the devices can be considered tolerant.

Given the likelihood that a traditional top-down systems engineering method will not suffice in addressing all issues encountered in low cost satellites based on commercial electronics, an alternative systems engineering approach was coined in (Engelen, et al., 2011), designed specifically for satellite swarms, which aims at incorporating emergent behaviour into the overall system design.

The method is outlined in Fig. 3.1, and starts by defining the desired global system behaviour. An initial estimate of the global swarm design, based on the best estimate of the design (and more precisely, the practical limitations, which are again fed back into the global swarm design) of the individual elements present in the swarm is then created. From this global swarm design, behavioural rules are created for the individual elements, which are then fed into the design of the individual elements. Through a process of integration, the resulting element design is then simulated into a swarm of identical elements, which is then verified against the desired global swarm design. Once a satisfactory solution has been converged to, the individual element behavioural rules can be frozen, as well as the global swarm design. Detailed design of the swarm elements can then follow, which should retain the behavioural rules dictated in this part of the design

process. An initial satellite swarm is then designed based on this desired behaviour, through decomposing the global swarm design into behavioural rules for each of the elements, which are then designed, based on this behaviour. Since the element interactions will cause emergent behaviour, system-wide simulations then provide feedback on how well the initial behavioural and functional objectives were met.

Swarm designs therefore are effectively iterated with varying element designs and element behavioural rules. Those rules define the behaviour of the satellite, and are in fact the primary target for the swarm design, as this allows limiting or tuning the influence and the effects of emergent behaviour. Note that, in contrast to traditional satellite design practices, satellite swarm design should specifically focus on the *behaviour* of the swarm, instead of purely on *functionality*, as the functionality of the individual elements, as single satellites, can be verified separately using traditional engineering methods.




Fig. 3.1: The proposed systems engineering method for satellite swarms (Engelen, et al., 2011)

The reason for the alternative systems engineering method is that the global behaviour of a swarm, which in effect is a virtual object, is defined by the behaviour of the individual elements, and more importantly, their interaction with each other. The method therefore also heavily depends on the actual element design, which can have its own distinct effects on the global swarm behaviour.

It is important to note that it is likely that not all emergent behaviour can be predicted or even simulated. A certain degree of safety measures, such as communication safeguards like message time-outs, and the ability to (temporarily) disable certain faulty elements, are recommended. Progress has been made in recent years however to predicting and managing the emergent behaviour resulting from a given element design (Winfield, et al., 2005), (Winfield, et al., 2005).

3.1 SIZING OF A SWARM

The target number of elements in a satellite swarm, as defined in (Engelen, et al., 2011) can be adapted to:

The number of elements, and their absolute or relative orbital position should be such that at least one (available) element is present at a given time slot at a given position in space.

This definition defines, based on a given observability of a certain phenomenon, the minimum number of elements required in the swarm. It assumes a perfect distribution, 100% reliability and permanent availability. This definition holds for any distributed system, and is designed to ensure timeliness of an observation. In an extreme case, one satellite will occupy an orbital slot as large as the satellite's observational area, which would imply that the entire orbit would be filled with equidistantly spaced satellites. It should go without saying that the timeliness requirement can be tailored to allow for a more realistic or more cost-effective scenario.

In reality, swarm satellites are not necessarily able to comply with this statement (Engelen, et al., 2011), due to satellites experiencing upsets, imperfect orbital distributions, or failed or degraded satellites, to which end a certain degree of over-definition of the number of elements is required. This has been treated in detail in (Engelen, et al., 2013), as well as in (Engelen, et al., 2012), and will be discussed further in the following sections.

3.1.1 The effects of numbers

Satellite swarms involve large numbers of identical satellites. Their large numbers enable them to increase revisit times in case of observation missions, and potentially allow using other satellites as hot spares of each other; taking over the tasks of another satellite in case it fails. With traditional satellites, the availability of the satellite, as well as the overall reliability is sufficient to ensure the satellite is capable of, and available to perform a given observation. Swarm satellites which have adopted a simplified design in order to save cost, which in turn allows for larger numbers of them to be launched within the same overall budget, may be less reliable than traditional satellites. Also, these satellites may be unavailable at times, which can interfere with an observation. They are therefore also dependent on the fact a larger number of them is launched into orbit, as otherwise certain mission goals would not be achieved.

The effect of numbers therefore plays a role in three distinct areas: overall swarm reliability, availability and coverage. When deciding on the number of satellites in a swarm, the required operational lifetime of the system (L), the effective lifetime of the individual satellites (l), and the minimum required number of operational satellites (m) in order to achieve the mission goals are important.

In this case, the required operational lifetime is the time the system has to remain operational, with at least m satellites remaining. The effective lifetime of the satellites l , is defined as the time the satellite is able to perform its nominal operations.

Given a known probability distribution of the individual satellite effective lifetime, it becomes possible to determine the minimal required number of satellites in the swarm,

for a given mission duration, when modelling the swarm as a k-out-of-m system. This model will be treated in more depth later in this section.

The following example uses a Gaussian distribution for the chance of a satellite failing as a simplified demonstrative case. Consider m to be the minimum required number of operational satellites, and k be the number of satellites present in the swarm. Then:

- In case the minimum expected lifetime of the individual satellites exceeds the required operational duration of the mission, the system can be designed with $k = m$ (Fig. 3.2, intersection (a)), as only aspects such as coverage, throughput and revisit times are driving the swarm design, as determined in Eq. (3.1).

$$k = m \quad \forall (l \geq L) \tag{3.1}$$

- In case the minimum expected lifetime of the individual satellites is less than the required operational duration of the mission ($l < L$) (Fig. 3.2, intersections (b) and (c)) however, the probability distribution will define how many satellites remain after a given amount of time, and the number of satellites in the swarm will have to be increased to account for satellites failing prior to the end of the mission. For a Gaussian distribution, this is shown in Fig. 3.2.




Fig. 3.2: Hypothetical probability distribution of failure of individual satellites in a satellite swarm, with three distinct cases highlighted.

For almost all probability distributions, there will be a number of satellites ' α ' after which adding extra satellites at the start of the mission to a swarm has minor effects on the overall lifetime of the swarm. A swarm with a total of $n = \alpha + m$ satellites at launch therefore represents an optimal swarm, optimised for system cost (Fig. 3.2, intersection (b)), as adding more satellites will increase the cost of the mission, without a noticeable effect on the overall lifetime (Fig. 3.2, intersection (c)). Note that cost is assumed to increase linearly with each added satellite, which inherently ignores the effects of mass production on cost, and that replenishing a swarm during the mission could result in a

dramatic cost reduction, as actual in-flight performance can be gathered, increasing the accuracy of the estimation of the required number of satellites.

It should be noted that infant mortalities are not taken into account in this scenario. Infant mortality would lead to a reduction in the number of available satellites at, or close to the start of the mission. In case an estimate of the infant mortality rate is available, this rate can be translated into a number of extra satellites which will have to be added to the swarm in order to overcome the expected infant mortality number. This procedure then finally renders the number of satellites required to guarantee a certain mission lifetime.

The number defined above satellites however does not necessarily have to be sufficient to guarantee coverage, as it only defines that the satellites will endure for the mission duration. In case of satellite clouds for example, none of the nodes features any significant form of orbit maintenance. Especially for such a system, additional satellites will be required to ensure that the desired target areas are covered by the cloud. However, at a certain point, the cost difference between the cheaper satellites of a satellite cloud, and those of a foraging satellite swarm for example would turn the scales towards a foraging swarm, which can ensure an even distribution of above the target areas in question.

A swarm can also be modelled as a parallel k -out-of- m system, as shown in Fig. 3.3 (Engelen, et al., 2013). In this model, m out of the k satellites present in the swarm are required to remain operational in order for the swarm to remain effective.

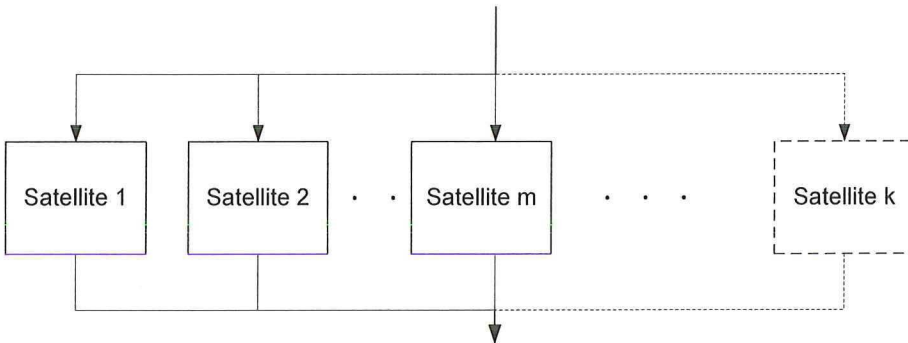


Fig. 3.3: k -out-of- m model of a satellite swarm

Given an estimate of the operational lifetime of an individual satellite, this model can be used to determine the effect of increasing the number of satellites in a swarm.

The lifetime, or Mean Time To Failure (MTTF) of this model can be computed through Eq. (3.2), in which s is the complex number frequency and λ the failure rate of a single satellite

$$MTTF = \lim_{s \rightarrow 0} \left(\frac{1}{s + (k - m)\lambda} \right) \quad (3.2)$$

or equivalently:

$$MTTF = \frac{1}{(k - m)\lambda} \quad (3.3)$$

The difference between k and m can be also be expressed as a number of extra elements in the swarm, as shown in Eq. (3.4), in which e represents the number of extra elements.

$$e = (k - m); (e = \alpha) \quad (3.4)$$

Using this model, the relative increase in lifetime, resulting from the presence of extra satellites in the swarm, can be computed. This is shown in Fig. 3.4, which shows the increase in lifetime of a satellite swarm consisting of up to nine extra satellites, with a requirement for one active satellite. As the figure shows, adding a second satellite results in the largest net effect. This is shown separately as “gain” in terms of improvement in lifetime compared to a system with one less satellite. Adding more satellites could be a viable option to increase the lifetime of the system, yet two important remarks are in order:

- The Markov model assumes an exponential probability distribution of the system’s failure rates. Satellites and other complex systems have been shown to exhibit other failure distributions however ((Castet & Saleh, 2009), (Dubos, et al., 2009), (Nolan & Heap, 1978)), hence the exponential distribution will result in a conservative estimate.
- The cost models for the production and operation of satellites is not yet tuned to swarms of satellites, hence the benefit of an increase in the number of satellites will not necessarily result in a dramatic drop in the per-unit costs. This implies that adding more satellites will increase the cost of the system, which potentially exceeds the increase in system lifetime gained through the addition of the extra satellites.




Fig. 3.4: Lifetime improvement prediction for an $e = 9$ satellite swarm. Gain is calculated as the relative lifetime improvement with respect to the previous added satellite.

3.1.2 Swarm spacecraft lifetimes and reliabilities

Swarm satellites ought to be designed with low unit cost in mind. This will undoubtedly force designers to make use of lower cost commercial components and parts, which can prove to be less reliable when operating in the space environment. This in turn can cause

the spacecraft to be less reliable than traditional satellites. It therefore appears that the desired reliability will have to be provided through the large number of satellites, rather than the individual satellite lifetimes. However, in order to allow assessing the lifetime of a swarm of satellites, some notion on the reliability of the individual satellites is required.

One can model an individual swarm satellite as a Markov chain¹² of connected systems, each of which is assumed to fail after a certain amount of time. Certain systems, such as a local data-storage device, can be allowed to fail in succession, before the satellite is considered lost. This will result in a certain time-to-failure, which allows for identification of the sensitive components in the satellite, as well as for an analysis of the estimated time-of-operation of an individual swarm satellite.

A generic, centralised, swarm satellite can be modelled as shown in Fig. 3.5. In this model, the satellite is controlled through a central On-Board Computer (OBC), which controls the data flows from the instruments and the radios, as well as controls the operational states of each of the subsystems. A Power Supply Unit (PSU) independently powers the satellite. Further components include an Attitude Determination and Control System (ADCS), as well as an inter-satellite and data downlink. The satellite is completed by inclusion of a payload and a data storage unit. Soft errors are included as a separate entity, allowing for analysis of their effect on the overall reliability, as well as the availability of swarm satellites.

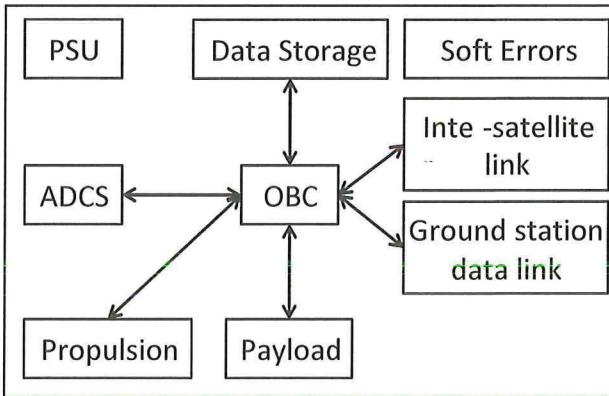


Fig. 3.5: Data-centric model of the swarm satellite as used in the Markov model

Since satellite swarms are decentralised by nature however, a centralised swarm node layout could prove to be less than ideal. Direct communication between the transceivers and the payload for example could be possible, with the OBC used primarily for orbit and swarm management tasks, as well as scheduling observations. This is depicted in Fig. 3.6. In this model, the payload directly places data into the data storage unit. The inter-satellite link and downlink have direct access to the data storage unit. The OBC schedules orbit correction and ADCS activities, and can control which of the payload data the data storage unit stores. Alternatively, the OBC can determine when to enable or disable the payload.

¹² Refer to Appx. A for a brief introduction into Markov modelling.




Fig. 3.6: Data-centric model of a decentralised swarm satellite

Using these models, it is possible to generate their respective Markov chains, highlighting the individual states the satellite is allowed to operate in. The chain based on the centralised satellite model shown in Fig. 3.5 is shown in Fig. 3.7. State “0” represents the state in which the satellite is fully operational. The subsequent states represent a certain state in which the satellite is operating with reduced functionality, due to the failure of a certain component.

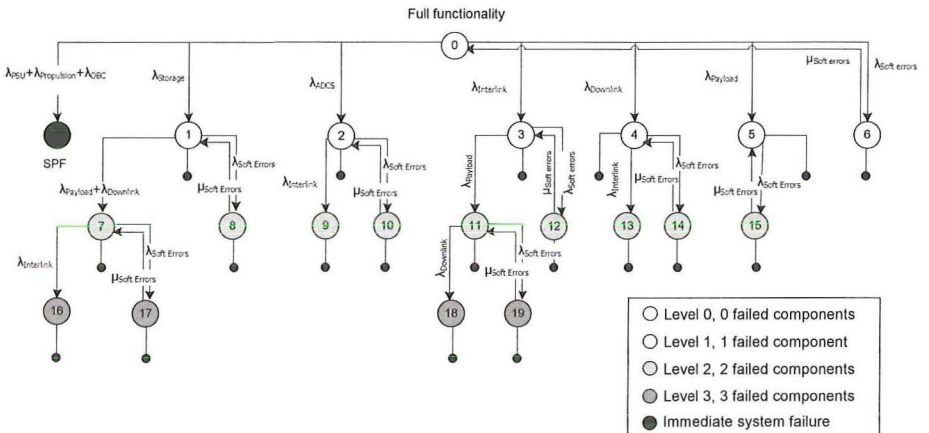


Fig. 3.7: The Markov chain for the centralised swarm satellite model

In this model, λ represents the failure rate of a certain component, whilst μ represents the repair-rate. Repairs in this sense consist of scrubbing of memories, or planned or unplanned resets of a digital system (in this case the soft errors). This is shown to effectively repair errors caused by single-event upsets (Carmichael & Brinkley, 2006), yet such procedures can result in a reduced availability of the satellite, which can be considered to be unavailable during such a procedure. This particular satellite has three Single Points of Failure (SPF), which are the PSU, the propulsion unit and the central OBC. With a broken OBC, no control of the satellite is possible, and no data will be stored, captured or transmitted. The propulsion system would result in an uncontrollable satellite as well, which would cause it to become a hazard to the rest of the swarm, which would have to move away from the dysfunctional satellite, to prevent collisions.

The storage unit is allowed to fail, in which case the satellite acts either as a relay station, or as a direct sampling device, only relaying recently captured data. The attitude control system is also allowed to fail, even though it is likely that high data rate communications would be disabled as a result. The inter-satellite link is also allowed to fail, in which case the satellite is assumed to start capturing data autonomously, using the down-link to relay it to the ground station. Lastly, the downlink of a certain satellite is allowed to fail, in which case, other satellites in the swarm are assumed to gather the data of that particular satellite through their inter-satellite link. Soft errors, which are a collective term for SEU-related issues, are assumed to be repaired, both pre-emptively, through resetting the system back to a known state at regular intervals, as well as through watchdog-systems protecting the satellite from system-hangs.

After one of these particular systems has failed, certain scenarios are modelled in which one or two more successive subsystem failures are allowed, as shown in Fig. 3.7. One example would be the successive failure of the inter-satellite link, followed by the payload. In that case, the satellite is still considered operational, even though its functionality is extremely limited. Should the downlink fail next, the satellite is considered lost, as it has no means of communicating, and will present a hazard to the swarm. This particular case should therefore be used to, once the inter-satellite link as well as the payload has broken down, remove the satellite from the swarm. In fact, the satellite must perform this action autonomously, in order to avoid having the ground segment instruct all other satellites to avoid the last-known position of the defective satellite.

A similar model can be created for the decentralised satellite model shown in Fig. 3.6. The resulting model is a lot larger, with the number of partially operable states increasing from 19 to 41, as many more failures can occur without the satellite failing in its entirety. The main advantage is that the OBC does not act as a single point of failure anymore, as can be seen in the model, shown in Fig. 3.8 and Fig. 3.9.

Due to the decentralised nature of the model, significantly more states (43 in total) are present in the model. In fact, the model was limited to three successive failures, even though more successive failures would be imaginable. As Fig. 3.8 shows, the decentralised model maintains two single points of failure, being the power supply, and the propulsion system. The notion is that satellites without a propulsion system cannot maintain their orbit and form a threat to the safety of the swarm. The satellite is then assumed to signal the failure, allowing the rest of the swarm to move away from the defunct satellite. Each of the other systems aboard a satellite is allowed to fail. Contrary to the centralised satellite, the OBC is allowed to fail. However, if the OBC has failed, both the attitude control system and the propulsion system are inaccessible, as the swarm control is assumed to reside in this computer. This is in direct conflict with the assumption that the propulsion system forms a single point of failure whilst the OBC does not, though scenarios in which the propulsion system performs a continuous thrust manoeuvre to move away from the swarm in case of a failure in the OBC are imaginable as a countermeasure. Those would not be available in case of a failed propulsion system.

The distinction between a failed OBC and any other system results in two state sinks: one in which the propulsion system still plays a role as a single point of failure mechanism, and one in which it doesn't anymore, as the OBC cannot control it anymore.

Fig. 3.9 then shows the remaining systems, which are the inter-satellite link, the downlink and the payload. Soft-errors are included in each of the branches, directly in the single-

point-of-failure state-dumps. Soft errors are highlighted only in a level 1-scenario, as a repair procedure can return the system into a fully functional satellite at this point.

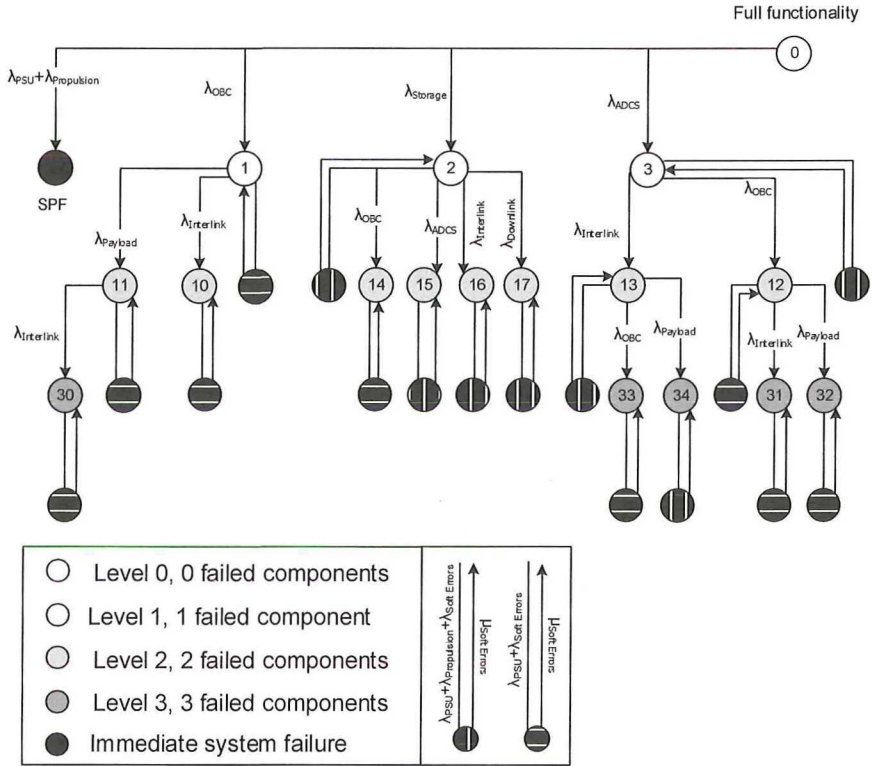


Fig. 3.8: The Markov chain for the decentralised satellite model (a)

Since no such satellite has been launched, or even been designed, to date, no information on the reliability of such satellites is available. Reliability data on small satellites has been gathered and analysed however (Dubos, et al., 2009). Swarm satellites are simple satellites, featuring virtually no redundancy in components, and, in order to save cost, are assumed to apply mainly COTS, mainstream components. They therefore, out of all satellite types launched to date, most closely resemble nano-satellites when considering their internal structure, especially given that nano-satellites have adopted a significant degree of standardisation, which facilitates production in larger numbers. Data on nano-satellites is available as well (Guo, et al., 2014), (Monas, et al., 2012), even though the number of launches, and subsequently the quality of data is significantly less. It is assumed the data on nano-satellite reliability will improve over time however, as more data becomes available.

Both Monas and Dubos report reliability figures for small satellites using a Weibull reliability distribution. A Weibull reliability distribution is defined through

$$R(t) = e^{(-t/\eta)^\beta} \tag{3.5}$$

In this case, $R(t)$ represent the reliability over time, and β and η represent scale parameters. In case $\beta = 1$, this equation is equal to an exponential distribution, and in case $\beta < 1$, components do not fail due to wear-out phenomena but are dominated by infant-mortality. This is true for most satellite components and missions, as well as other complex systems, built in low production volumes (Nolan & Heap, 1978). Incidentally, $\eta = 1/\lambda$, which allows using the scale parameter given for a Weibull distribution for the failure rate required for Markov chain reliability analysis.

When doing so, failure rates for the various components in a swarm satellite can be derived, as shown in TABLE 3-I.

TABLE 3-I
ASSUMED SUBSYSTEM FAILURE RATES,
ADAPTED FROM (MONAS, ET AL., 2012) AND (CASTET & SALEH, 2010)

Failure rate		
Subsystem	Nano-satellite	Traditional satellite
PSU	1/229 years	1/169272 years
OBC	1/2712 years	1/7983 years
ADCS	1/455 years	1/3831 years
Storage	1/2712 years	1/7983 years
Inter-satellite link	1/814.1 years	N/A
Down-link	1/814.1 years	1/400982 years
Payload	1/2712 years	1/7983 years
Propulsion	N/A	1/6206945 years
Soft error rate	1/10 hours	N/A

These failure rates can be used to determine the Mean Time To First Failure (MTTFF) and the Mean Time To Failure (MTTF), as well as the point in time at which the reliability drops to a given value. The reliability point was set at 90% for the calculations performed in (Engelen, et al., 2014), using the Markov models shown in Fig. 3.6, Fig. 3.7 and Fig. 3.9, and the reliability point was calculated from the time to first failure.

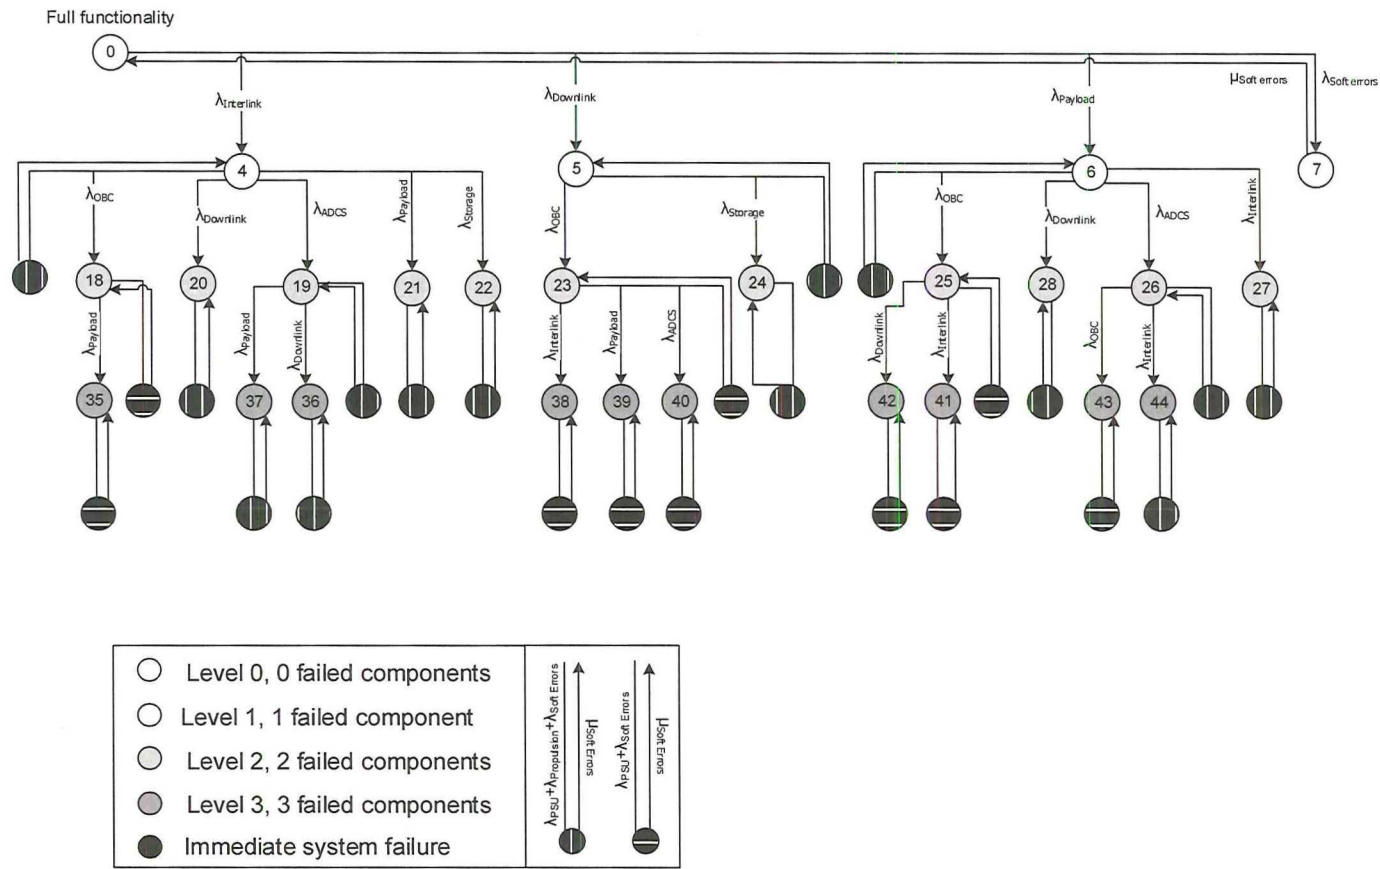


Fig. 3.9: The Markov chain for the decentralised satellite model (b) continued from Fig. 3.8.

Three distinct scenarios were computed, for both the centralised and decentralised satellite models:

1. A reference scenario, in which single event upsets (represented through the soft-error branches) were not considered
2. A scenario in which soft-errors occurred, yet no countermeasures were taken
3. A scenario in which soft-errors occurred, against which the satellite pre-emptively reset all systems at a rate at which the satellite remained available for 95% of the time.
4. A scenario in which soft-errors occurred, and the satellite pre-emptively scrubbed all systems, yet at an optimised rate¹³.

The results are summarised in TABLE 3-II. For computing these results the swarm satellite failure rates were taken equal to those of a nano-satellite.

TABLE 3-II
CALCULATED MEAN TIME TO FAILURE AND MEAN TIME TO FIRST FAILURE OF
INDIVIDUAL SWARM SATELLITES

Scenario	MTTF ^A	MTTFF ^A	90% reliability point
Centralised satellite model, excluding soft errors	119 years	81 years	8.54 years
Centralised satellite model, including soft errors, excluding countermeasures	20 hours	10 hours	1 hour
Centralised satellite model, including soft errors. Repair rate such that the satellite is available 95% of the time	100.5 years	80 years	8.39 years
Centralised satellite model, including soft errors, optimised repair rate	104 years	73 years	7.68 years
Decentralised satellite model, excluding soft errors	9 years	3.3 years	0.36 years
Decentralised satellite model, including soft errors, excluding countermeasures	20 hours	10 hours	1 hour
Decentralised satellite model, including soft errors, repair rate such that the satellite is available 95% of the time	42 years	3.3 years	0.35 years
Decentralised satellite model, including soft errors, optimised repair rate	16.5 years	3 years	0.32 years

^A Calculated using input data as given in TABLE 3-I, for component failure rates taken equal to those of nano-satellites.

As the table shows, the MTTF of the centralised satellite model exceed those of the decentralised satellite model. In case soft-errors are present, yet no countermeasures are taken, both satellite models encounter their first error, as expected, at the occurrence of the first soft error. Since the soft-error branch is only one state deep, the second occurrence results in the system entering a Single Point of Failure state, which implies the satellite is lost. For the centralised model, the MTTFF and MTTF are quite close, yet

¹³ Optimising the repair rate is done through matching the repair rate with the expected upset rate (see (Engelen, et al., 2014)).

for the decentralised model, these times lie much further apart. This is due to the larger number of states present in the decentralised model.

It becomes immediately apparent that the decentralised satellite model is less sensitive to soft-errors though. In fact, the repair operation appears to increase the satellite MTTF to a time beyond the reference case. This is due to the cross-links in the model, in which the model can remain in a state with operational OBC for much longer. However, the MTTF is a much more relevant figure, or in case of single satellites, the 90% reliability point as at that point, the satellite has a 10% chance of having failed already. In the case of satellite swarms, this has less of an impact on the overall system, yet for single satellites, this is obviously quite a high risk already.

Interestingly, the computed lifetimes for the centralised model are significantly longer than the ones observed for actual nano-satellites, which are reported to reach their 90% reliability point already in the first six months after launch (Guo, et al., 2014). This is mainly attributed to the high number of infant mortality cases, which appear to have been taken into the analysis. In fact, the computed lifetimes for the centralised model appear to be closer to those observed in small satellites (Castet & Saleh, 2009). The discrepancy lies mainly in the low reported failure rate of the OBC, which in turn could be improved with more data on actual flight experiences of nano-satellites. Also, with more nano-satellite launches every year, more experience will be gained with time passing, likely causing a shift in the expected lifetime of such satellites as well.

The decentralised satellite model results in initial lifetimes which lie much closer to the MTTF's reported in literature, upon which the input data for these models were based.

The computed times to failure however are much shorter than those computed for the centralised model, which seems counter-intuitive, as the decentralised model ought to be more robust against component failures. The root cause lies in the size of the resulting model: when the number of states increases, the chance of remaining in the '0'-state, which represents the MTTF, reduces. This can easily be verified by using a simplified model, in which all failure rates are equal. The MTTF and MTTF can then be represented by Eq. (3.6), in which n represents the number of states present in the model

$$\begin{aligned}
 MTTF &= \lim_{s \rightarrow 0} \left(\frac{1}{s + n\lambda} \right) \\
 MTTF &= \lim_{s \rightarrow 0} \left(\frac{1}{s + n\lambda} \left(1 + \frac{n\lambda}{s + n\lambda} \right) \right).
 \end{aligned}
 \tag{3.6}$$

The resulting MTTF, along with the time spent in the '0'-state (i.e. the MTTF), is shown in Fig. 3.10, for increasing numbers of states. As expected, increasing the number of states decreases the MTTF, as well as the MTTF. A direct comparison of the failure rates reported in these models is therefore not possible. Care should therefore be taken to only examine the relative changes.




Fig. 3.10: MTTF and MTTF for a simple system with an increasing number of states

Perhaps the most remarkable aspect of the decentralised satellite model is that the MTTF, when increasing the repair rate, exceeds the MTTF of the same satellite model which does not encounter single events, which would represent the ideal case. This can be attributed to the fact that for this model, the SPF-state, which acts as the state dump¹⁴ for a failed satellite, includes a repair mechanism for soft-errors. This implies that the satellite can emerge from the 'failed' states through a repair procedure. While this only holds for errors caused by single events, this is frequently observed even on Earth, in commercial applications such as mobile phones and desktop PC's, after a software crash. What is remarkable however is that, in this case, repairing more often might prove beneficial after all, as the resulting satellite availability will not suffer as much as in the centralised case, as in the decentralised case the OBC, which is the most likely to be affected, is a decentralised component.

Fig. 3.11 shows the time spent in each of the reduced operational states for the centralised satellite model, for each of the modelled scenarios. State 6, which represents the soft-error branch of the Markov tree clearly shows that repairing at the optimal rate significantly increases the time spent in that state. Also state 12, which represents the case where soft-errors occur after the inter-satellite link has failed shows a similar characteristic. Those two states combined then make up for the reduced time spent in the other states, compared to the SEU-less case; resulting in a reduced, yet still acceptable MTTF.

¹⁴ A state dump represents a state which terminates a branch, resulting in a failed satellite.

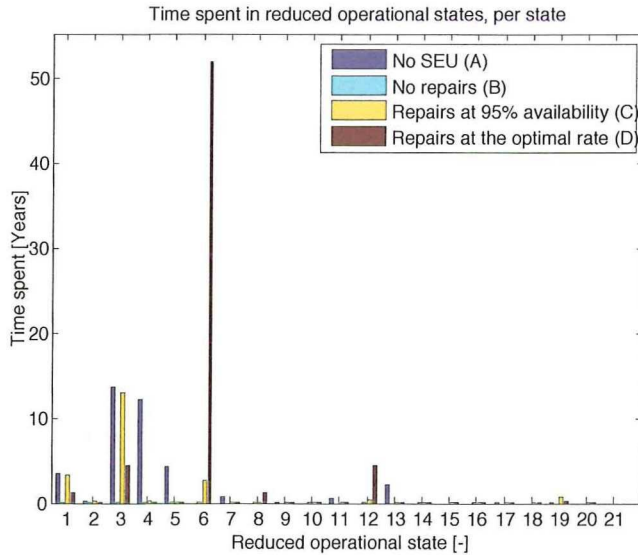


Fig. 3.11: Time spent in reduced operational states for the centralised satellite model.

Repairing more often, represented by case 'C' also increases the time spent in state 6, though significantly less than for the optimal repair rate case. Most other states then approach the times reported for the SEU-less case (case 'A'), which implies this model approaches the ideal case for those states.

For the decentralised model, the differences between each of the four different repair scenarios are more pronounced, both in terms of time available as a fully functional satellite; as in the overall MTTF. Fig. 3.12, Fig. 3.13 and Fig. 3.14 show the time spent in each of the various states, for each of the repair scenarios. The extremely high MTTF reported for the 95% availability scenario can easily be explained from these figures, as the time spent in states 2, 4 to 5, 10 to 17, 35 to 39 and 50 to 53 is significantly longer for the 95% availability scenario than for any of the other scenarios. These states mainly involve the ADCS, the inter-satellite link, and the down-link as the involved systems, which incidentally are also the systems with the shortest times-between-failures. Repairing more often is therefore likely to reduce their influence on the overall (system) failure rate, which would cause the increased time spent in these states.

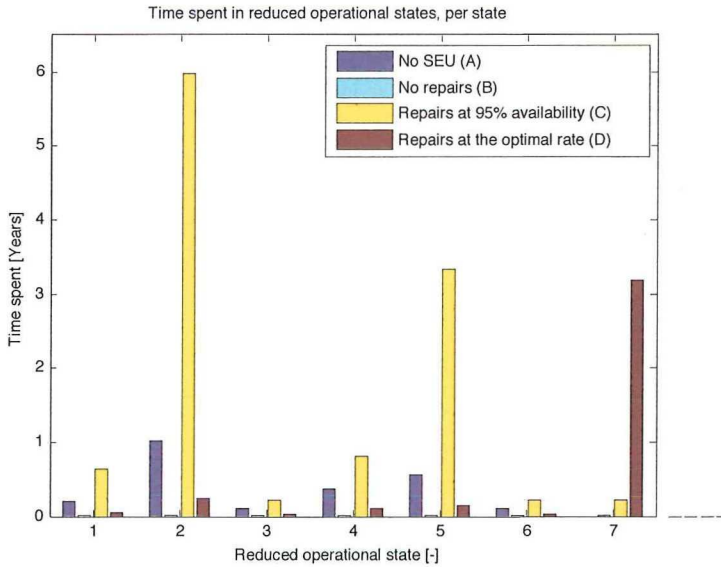


Fig. 3.12: Time spent in the level 1 reduced operational states for the decentralised satellite model.

For the optimal repair rate case, the time spent in state 7 is, as expected, significantly longer, at slightly over three years in total. This is entirely in line with the centralised satellite model, shown in Fig. 3.11, in which this branch of the Markov tree is represented by state number 6. All other states are rather uneventful, save perhaps for states 29-34, which in total amount to less than 1.3 years. The fact the repair-less case is invisible in the graph above can be explained, in that as the total MTTF is only 20 hours none of the states are visible on the time scales shown in the graphs.

The cost aspect of designing a satellite swarm should not be neglected. In certain cases, a more expensive yet more reliable satellite will reduce the number of satellites required in order to meet the mission duration criteria. If this reduced number of satellites is then also able to meet the observability criterion, this swarm based on more expensive spacecraft could prove to be more economical overall. The analysis presented above can be repeated for different satellite designs, with different associated (predicted) costs, after which comparative conclusions as to which design is more cost effective can be made, serving as useful input into the mission and system design process.

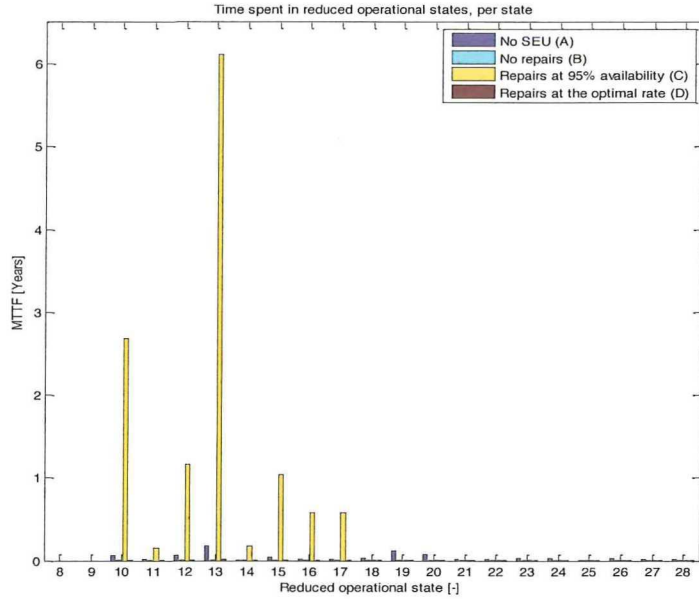


Fig. 3.13: Time spent in level 2 of the reduced operational states for the decentralised satellite model.

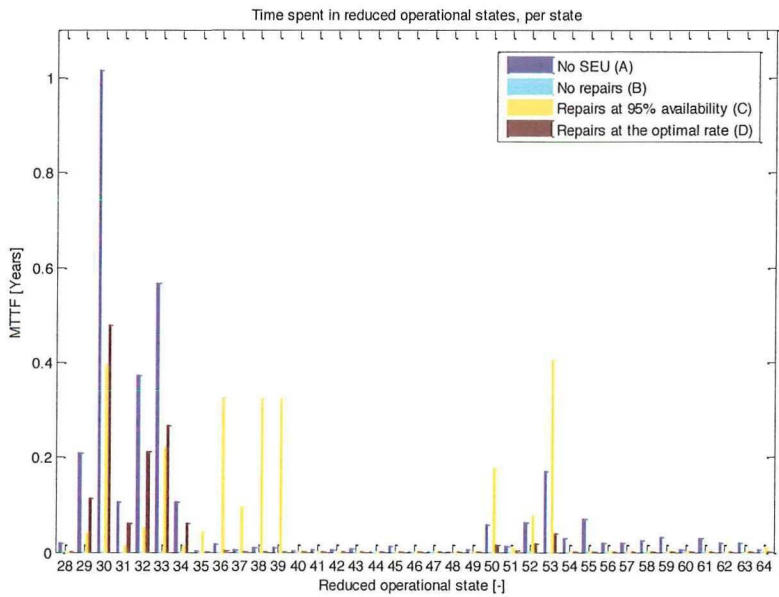


Fig. 3.14: Time spent in level 3 of the reduced operational states for the decentralised satellite model.

3.1.3 Overall swarm lifetime

Swarms are said to allow for increased system lifetimes, which can be explained through the inherent redundancy offered by the large number of spacecraft (Verhoeven, et al., 2011).

In terms of cost however, increasing the number of satellites in the swarm to increase the lifetime could prove to be prohibitively expensive. It would therefore be useful to obtain an accurate estimate of the lifetime of a swarm based on a given element design, in order to assess the suitability of that particular element design in achieving the intended mission duration.

Determining the lifetime (or the MTTF) of a particular swarm can be done using the k-out-of-m system model described in section 3.1.1, using the element lifetimes computed in section 3.1.2. In (Engelen, et al., 2013) however, a more in-depth Monte Carlo analysis was performed, using a Weibull distribution for the satellite failure rates, rather than the exponential distribution which is assumed in case of a Markov chain. A Weibull distribution would match the observed failure rates, reported in (Monas, et al., 2012) more closely, resulting in a more realistic system lifetime estimation.

The parameters used for the Monte Carlo simulation are shown in TABLE 3-III. As can be seen, the downlink system was assumed to have an excess capacity of 10%, and the swarm was assumed to consist of 100 satellites. The individual satellite lifetimes were spread according to the Weibull distribution for each of the 100 satellites, and the numerical average was taken over each different run.

Spreading the satellite lifetimes was performed using a random number, picked from a Weibull distribution. This can be achieved through applying Eq. (3.7) to a random number x , selected from a uniform distribution

$$y = [-\eta \ln(x)]^{\frac{1}{\beta}}, \forall x \in [0..1]. \quad (3.7)$$

In this equation, β and η represent scale parameters, as was the case in Eq. (3.5), and it can easily be derived through solving Eq. (3.5) for t , which in this case is represented by the random number x .

Subsequently, the lifetime of each individual satellite, in case no manufacturing deficiencies are taken into account, is scaled according to Eq. (3.8), which results in a lifetime of each of the satellites which is *at most* equal to the lifetime computed in section 3.1.2,

$$l_{sat} = \left(1 - \frac{y}{\max(y)}\right) \frac{1}{\lambda_{sat}}. \quad (3.8)$$

Note that the lifetimes computed in section 3.1.2. assume an exponential distribution, which is conservative compared to the observed satellite lifetimes, which are shown to exhibit Weibull distributions. The Monte Carlo method, which uses the computed lifetimes as input parameters, will therefore render conservative overall swarm lifetime estimates.

TABLE 3-III
INPUT PARAMETERS TO THE MONTE CARLO ANALYSIS

Parameter	Value
Lifetime variation distribution	Weibull
Shape parameter	0.3134 ^(A)
Number of satellites in the swarm	100
Spare bandwidth capacity per element	10%
Manufacturing deficiencies modelled as (where applicable)	Normal distribution, 40% spread, 3σ

^A Data taken from (Monas, et al., 2012)

As reported in (Engelen, et al., 2014), and repeated here in Fig. 3.15, the lifetimes predicted by the Monte Carlo simulation, which in this case are represented by the Mean Time To Failure (MTTF) of the satellite models, show that most satellites remain operational for at least 100 years. Once a certain threshold has been reached however, which appears to lie around 90 remaining satellites, the whole system appears to start failing dramatically.

The Monte Carlo simulation, like the satellite lifetime determination detailed in section 3.1.2 assumes four distinct cases. In case A, no Single Event Upsets (SEU) are assumed to occur which coincides with no soft errors occurring. This case represents the theoretically achievable lifetime. In case B, SEU's occur, but no actions are taken against them, which results in a total satellite lifetime of only 20 hours. This case is not shown in Fig. 3.15, as it would not be visible. Case C and D then consider satellites encountering soft errors, yet countermeasures are taken. The rate at which these measures are taken is varied. In case C, the satellites repair such that they retain an availability of 95%, and in case D, the repair procedures are done such that the satellite maximises availability, at the expense of a shorter overall lifetime.




Fig. 3.15: Monte Carlo simulation result for the lifetime of a 100-satellite swarm for the centralised satellite model, without taking manufacturing deficiencies into account.

This simulation can be repeated for the decentralised swarm satellite model, which results in the simulation result shown in Fig. 3.16, for 100,000 runs, taking manufacturing

deficiencies into account. As the figure shows, the lifetime of the swarm almost equals the lifetime of an individual element, plus the assumed spread due to manufacturing deficiencies. Infant mortality cases were not taken into account in this simulation, hence the assumption that all 100 satellites last until their designated lifetime.

The resulting lifetimes of the individual satellites are much longer than the observed lifetimes of nano-satellites, of which the failure distribution data were used as input parameters. This then also results in a very long swarm (system) lifetime. The cause lies primarily in the Markov model's simplicity. Simple satellites have few (connected) failure cases, which gives them a long theoretical lifetime. The models used also allowed for continued operation with failed components, which increases the predicted lifetime of the satellite even more. Comparing the Mean Time To First Failure should therefore show a closer match to observed lifetimes of satellites, yet the simplicity of the centralised swarm satellite model still causes a significant discrepancy. The decentralised satellite model then shows MTTF's which lie much closer to the MTTF's observed in the nano-satellite missions on which the input data was based, which is mainly due to the increased number of states.

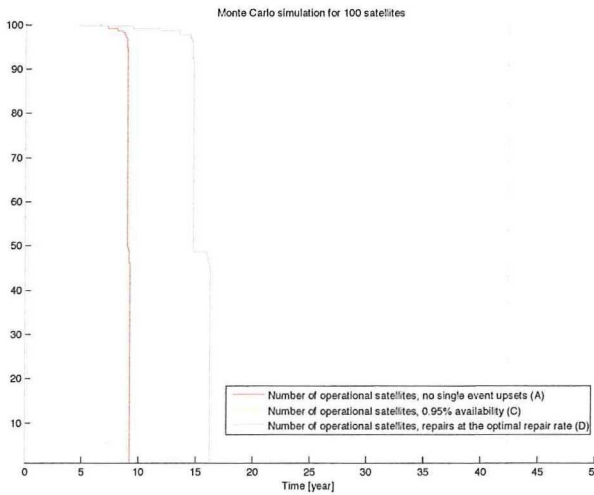


Fig. 3.16: Monte Carlo simulation result for the lifetime based on the MTTF of a 100-satellite swarm, for the decentralised satellite model, taking manufacturing deficiencies into account.

The most interesting aspect however lies in the difference between the lifetime of a single satellite, and the lifetime of the swarm. This is shown in Fig. 3.17, in which the theoretical lifetime of a single satellite is shown as a dashed line, and as can be seen, the swarm as a whole *can* reach the lifetime of, or even outlive the individual satellites, though with the used simulation parameters, not by much. The cause of this lies purely in the spread caused by manufacturing variances, which, when modelled as a normal distribution results in half of the satellites surviving for more than the nominal lifetime. This is the primary cause for the increased reliability commonly associated with swarms, as even unreliable elements can rely on the redundancy presented by the swarm. Note

that in all cases, even unreliable satellites should remain functional in one way or other, or at least be able to return to a functioning state in time.




Fig. 3.17: Monte Carlo simulation result showing the difference between the lifetime of a single (ideal) satellite, and the lifetime of the swarm

3.1.4 Availability and throughput

Satellite swarms, assuming sufficient inter-satellite bandwidths and downlink bandwidths, would in theory be capable of gathering and distributing tremendous amounts of data. However, given the assumption of cheap, yet slightly less reliable satellites, not all satellites are guaranteed to be operational. This could pose problems for certain missions which rely on rapid acquisition of time-critical data, such as early-warning systems. Yet for all missions, the unavailability of a certain satellite will cause a reduction in the system's throughput, and, by extension, also its output. This was investigated in (Engelen, et al., 2012) as well as in (Engelen, et al., 2013), in which an assumption was made that soft-errors could be "repaired". This is nowadays an accepted fact, and many components flown in spacecraft today actively use a variety of repair mechanisms (Martin & George, 2012), counteracting some of the effects of charged particles on their digital components.

In FPGA's for example, scrubbing can occur on an active device, which implies the system which includes the FPGA can remain active (and therefore the satellite will remain available) throughout the procedure. For certain methods, this has proven to result in an availability of such systems in excess of 99.999% (Martin & George, 2012). Processor-based systems which only use a single processor cannot perform such procedures while remaining active. They can however sporadically perform a reset which would result in the system rebooting. This causes the processor to revert to a well-defined state. Given that most low-cost satellites show a slight preference for making use of processor-based (sub-) systems (Sweeting, 1992), (Asenek, et al., 1997), (Ginosar, 2012), it was assumed that swarm satellites would be no different, and that they would therefore also apply microprocessors. They would therefore also require regular or at least sporadic resets (Lovellette, et al., 2002).

For the analysis of the effects on availability and system throughput, it was assumed that a repair procedure would happen at regularly scheduled intervals, and that the procedure would take the satellite offline for a period of one minute per operation. This is quite conservative, given the speed of commercial electronics at the time of writing. Even so, as shown in Fig. 3.18, with a failure rate of one upset every 10 hours (which exceeds the observed upset rate reported by (Lovellette, et al., 2002)), each satellite can increase its MTTF significantly when applying less than one repair procedure per hour.

While Fig. 3.18 was made for the centralised swarm satellite model, other systems will show similar behaviour, as the net effect of a repair procedure is that the state transition for the soft-error state is prevented, which effectively eliminates it from the Markov chain model. Repairing after the soft error has occurred can cause the system to revert to the fully operational state, which is typically what would happen after a watchdog reset or after external intervention. These occurrences nonetheless are not taken into account in the analysis of the centralised satellite model. Note that the MTTFF in this case goes down. Since the MTTFF goes up however, it implies the system spends much more time in state '0', which is where it is most useful to the mission.




Fig. 3.18: Distribution of the MTTFF and MTTFF (1) of the centralised swarm satellite model.

Fig. 3.19 shows the result for the decentralised satellite model. In this model, the effects are even more pronounced, as soft errors form part of the single-point of failure state, which acts as a state-dump for a failed satellite. Satellites can therefore be “revived” from a failed state, provided the failure was soft-error related. This is most prominent in the fact that the MTTF is increased significantly through increasing the repair rate, whilst the MTTFF levels off in a manner quite similar to the effects seen in the centralised satellite model. Reviving a satellite from the failed state incidentally mimics the effect of an external reset, which is the most likely means of repairs to be used in low cost satellites.

This figure also shows a steep increase in MTTFF, which is due to larger number of repair options present in the system, allowing the system to continue nominal operations more often.

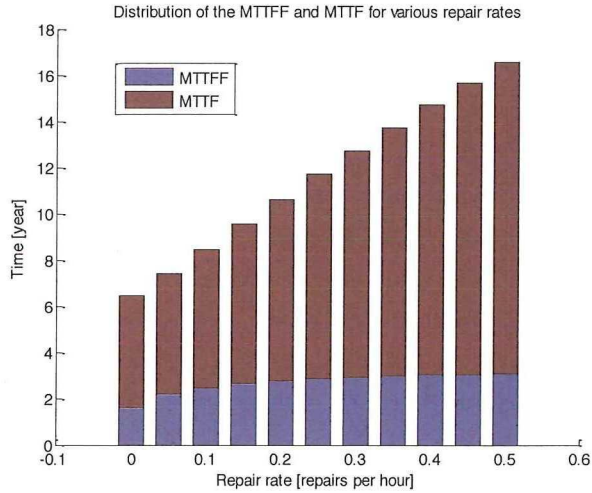


Fig. 3.19: Distribution of the MTTF and MTTFF (2) of the decentralised swarm satellite model.

Fig. 3.20 then shows the effect a repair procedure on the availability of the specific satellite, as well as the point at which the area under the curve is maximal. This would coincide with the point at which the repair rate has the maximal effect on the MTTF, with the least amount of impact on the availability (Engelen, et al., 2014). As can be seen, the availability can be optimised to 97.6%, when repairing at a rate of 0.4 repairs per hour. Increasing the repair-rate increases the MTTFF, at the expense of satellite availability. Moreover, increasing the repair-rate beyond one repair procedure per hour shows minimal effect on the MTTFF, yet the impact on the availability of the satellite drops significantly.

For the decentralised satellite model, the result is shown in Fig. 3.21, which shows an identical optimum in terms of repair rate, due to the identical soft error occurrence rate. The resulting MTTFF is reduced though.

The (data) throughput of a satellite swarm can be significant, due to the large number of satellites involved. Physical and legal limitations on the available bandwidth and use of frequencies will have an impact on actually achievable performance. Some of these were briefly discussed in (Rajan, et al., 2011), yet a more generic treatise of throughput was presented in (Engelen, et al., 2012). In this analysis, each satellite was assumed to have an excess bandwidth of 20% available which was to be used to replace the bandwidth of satellites with a broken down-link. Note, in this case all satellites were assumed to feature an inter-satellite link with sufficient bandwidth to distribute the data across all of the swarm and that each satellite has a downlink with sufficient bandwidth to download their obtained payload data.

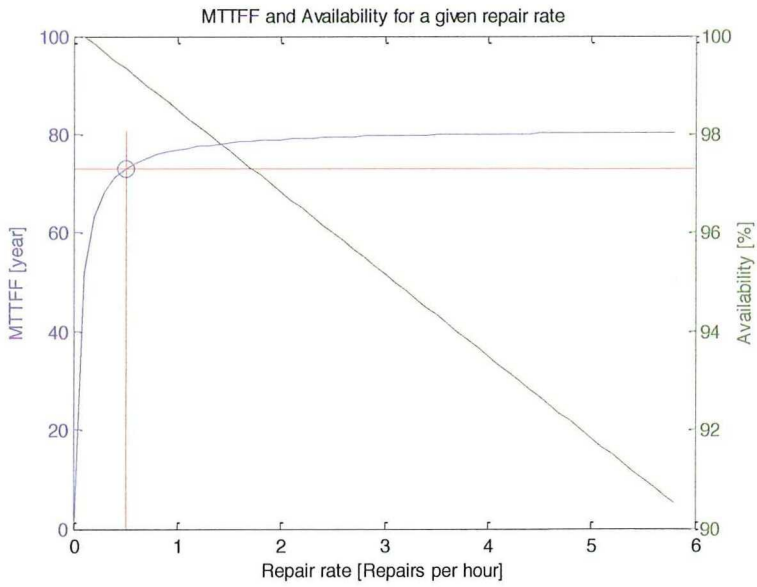


Fig. 3.20: MTTFF versus availability for the decentralised swarm satellite model. The circle indicates the optimised repair rate.




Fig. 3.21: MTTFF versus availability for the centralised swarm satellite model. The circle indicates the optimum repair rate.

The result of the analysis is shown in Fig. 3.22, which shows that the downlink capacity in terms of satellite units (the available bandwidth per functional satellite is '1.20') always exceeds the available payload data due to the fact that the aggregate system shows a shorter MTTF than the MTTF for the payload and the downlink. Given these assumptions, each satellite would be able to downlink all of the gathered data. Since the excess bandwidth per satellite, which in this case was assumed to be 20% of the total available bandwidth, also exceeds the reduction in their availability due to single-events and repair-procedures, each satellite can in most circumstances transfer all of its data to a ground station. The throughput of such a swarm would therefore be equal to the theoretical bandwidth required to downlink all of the gathered payload data. This analysis however does not include data lost due to single events occurring when storing acquired data. Those effects are assumed to remain limited. A worst-case estimate could be assumed, taking the availability of the satellites as the reference, in which case the average output of each of the satellites would be equal to the product of their availability times their bandwidth.

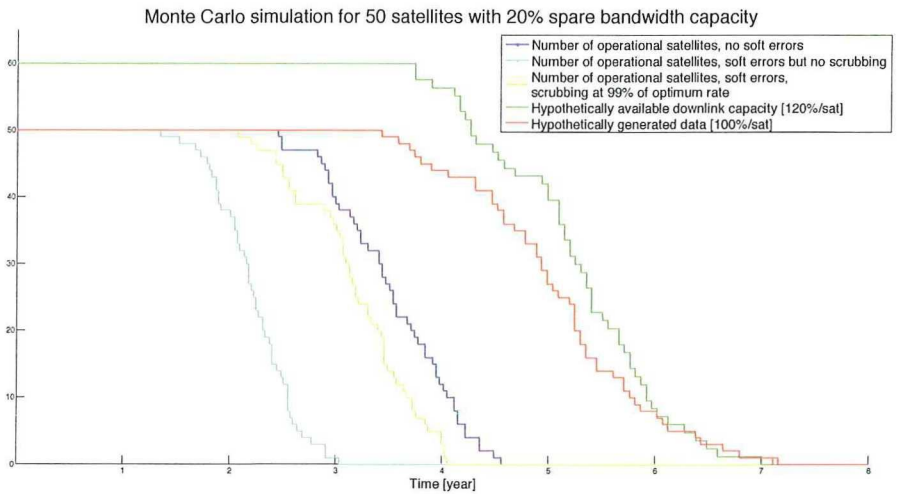


Fig. 3.22: Effect of failing satellites on the throughput of a 50-satellite swarm each with 20% spare bandwidth capacity.

3.1.5 Swarm system reliability

The reliability of a system, which assumes an exponential failure distribution can be defined as being equal to the probability that a system will last for a time equal to t :

$$R(t) = e^{-\lambda t}. \quad (3.9)$$

The reliability when assuming a Weibull probability distribution is defined by

$$R(t) = e^{-\left(\frac{t}{\eta}\right)^\beta}. \quad (3.10)$$

In case of a satellite swarm, the question arises whether the element reliability is still relevant, as the swarm acts as a redundant system. This was investigated for the centralised swarm satellite model in (Engelen, et al., 2013), and the result is shown in Fig. 3.23, which shows the result of the Monte Carlo simulations run for section 3.1.3 calculating the reliability of the system, assuming an exponential reliability distribution of the overall system.




Fig. 3.23: Reliability of a swarm over time using the centralised swarm satellite model

Fig. 3.24 then zooms in on reliabilities above 85%. What becomes apparent is that the single satellite reliability, determined using an exponential reliability model, exceeds the reliability of the swarm, which uses the Weibull distribution, after a period of about 25 years, whilst the swarm's reliability exceeds that of the single satellite during the period before. Traditional satellite systems are designed for reliabilities of 90% and above, which would imply that in all common cases, the reliability of a satellite swarm can exceed the reliability of the single satellite, in spite of its simplified hardware and reduced per-satellite reliability. In this case, the single satellite reaches a reliability of 90% after a period of 7.7 years, whilst the swarm reaches that level after around 11 years.

The margin by which the swarm reliability exceeds the individual element reliability ultimately depends on the satellite model used, as well as the failure distribution.




Fig. 3.24: Reliability of plot of the centralised swarm satellite zoomed in for reliabilities above 85%

The implications of the finding that swarm reliability can exceed that of an individual satellite are significant, and the conclusion is that satellite swarms should exploit their great numbers, not only for an increase in system reliability, but also for the resulting increase in throughput (Engelen, et al., 2014). As swarms are more reliable during the first part of their operational lifetime, they should gather as much data as possible during that phase of the mission; after which they could be replaced.

Satellite swarms, when designed similarly to the swarm modelled above, therefore do not have to rely as much on the reliability of the individual satellites. Other satellites can assume the tasks of a defective satellite, implying that designing swarm elements with high individual reliabilities will likely drive up the cost, whilst providing little additional benefit to the swarm.

A similar analysis can be performed for the decentralised satellite model. The result is shown in Fig. 3.25 and Fig. 3.27, for a Monte Carlo simulation run of 100,000 iterations. Two distinct versions are generated: Fig. 3.25 assumes an exponential reliability model, as was done for Fig. 3.23 and Fig. 3.24. Since the reliability model used for the Monte Carlo simulation assumes a Weibull probability distribution it will be more representative if such a distribution is used for generating a second version of the reliability figures. This can easily be achieved through applying Eq. (3.10) to the output of the Monte Carlo simulations, and the result is shown in Fig. 3.26. Fig. 3.27 then zooms in on reliabilities of 85%.

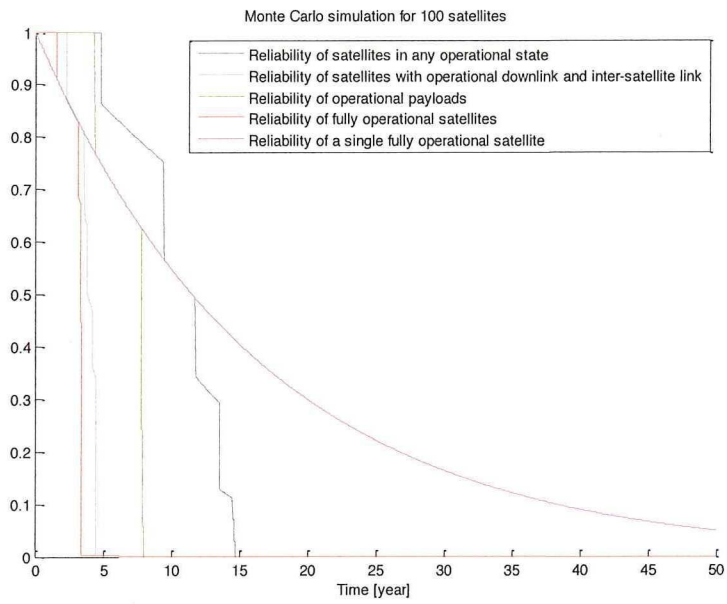


Fig. 3.25: Reliability of plot of the decentralised swarm satellite (1) for an exponential distribution.




Fig. 3.26: Reliability of plot of the decentralised swarm satellite (2) using a Weibull distribution




Fig. 3.27: Reliability of plot of the decentralised swarm satellite (3) using a Weibull distribution, zoomed in for reliabilities above 85%

As can be observed from these figures, the reliability of the swarm (shown as a black line for satellites in any (partially) operational state, and as a red line for fully functional satellites) exceeds the theoretically predicted reliability of a single satellite. However, at the point where the reliability of the swarm sinks below that of a single satellite, the reliability of the overall system plummets, implying that in case of satellite swarms with these properties attempts to use the satellites beyond their predicted lifetime should be discouraged. The optimal method would likely be to follow the reliability of the payloads, as satellites without a payload offer very little to the systems' throughput.

These results are in line with the expectation that adding more elements to a swarm would increase the reliability of the system (Verhoeven, et al., 2011), yet also show clear boundaries to the usefulness of such a method, as clearly shown by the plummeting reliability exhibited by the decentralised satellite model.

3.2 REDUNDANCY

As discussed in section 3.1.5, the reliability of a satellite swarm can exceed the reliability of the individual elements, provided that variations in either the manufacturing quality of the individual elements are present, or the effects of external influences vary across the swarm elements. This implies that the redundancy offered by the swarm is beneficial to the lifetime and system reliability in almost all realistic scenarios.

Traditional satellites however often apply redundant (internal) (sub-)systems, in order to increase their reliability and/or availability (Huang, et al., 2009). Swarm satellites are in fact redundant copies of one another, giving rise to the question whether swarm satellites should also feature internally redundant systems, as in the case of traditional satellites, or whether they can make due with a singular, non-redundant internal architecture. Singular internal substructures are quite common in commercial devices like mobile phones and other home appliances. In such systems, low cost is imperative, both during design, as well as during manufacturing. In case of a failure however, in-situ repair is available which is significantly different in the case of satellite electronics. Another, perhaps even more prominent feature is the very large numbers available for analyses, which in turn is used to tune manufacturing and design processes to render the largest yield and coincidentally also the largest reliability. Since satellites are essentially one-off products, produced in extremely small quantities per type, statistics generally do not apply, therefore very little feedback is available, preventing fine-tuning of manufacturing processes or system design practices (Tsinas & Welch, 2001). In terms of reliability, satellites require on-orbit validation which implies statistics only become available after a reported on-orbit failure. Given the long operational lifetimes of satellites, this is a process that can take years causing all statistical information to lag behind the state of the art by a significant amount. This is not likely to change in the near future though, as satellites are not likely to become a mainstream product. A large database (see for example (Castet & Saleh, 2009)) is maintained, analysed and published at regular intervals, reporting on the current state of affairs.

Internally redundant systems are commonly used to address random failures, which can be modelled as a parallel Markov chain, or as a k-out-of-m system. In case those are the dominant failure modes, the reliability of a system with internal redundancy can be shown to approach unity. This also holds for satellite swarms, as shown in section 3.1.13.1.4, which apply external redundancy in an analogous manner. In this case, the swarm as a whole serves as a model of an individual satellite, with each of the swarm elements as the satellite's internally redundant components.

One more pressing issue, as addressed by (Apostolakis, 1976) is that common mode failures¹⁵ are not addressed by internally redundant systems. He finds that common mode failures may dominate by a factor of 10^5 compared to random errors, and establishes that, for a system which does not allow for periodic inspection and in-situ repairs, common mode failures are the limiting factor for the lifetime of the system. In case repairs are available, an upper bound for the frequency of common mode failure occurrences can be found, when routine inspections are performed.

¹⁵ Common mode failures are failures with a common cause. For example, if a component is radiation tolerant to a certain total dose, all similar systems will fail at that total dose, and adding more identical systems will not help.

The question at hand is then which failure modes are most relevant in case of swarm satellites and what countermeasures would be available. Common cause failures in case of swarm satellites are limited to design flaws and manufacturing defects, as the satellites' physical separation ensures environmental impacts, such as damage caused by debris or damage caused by degradation, do not simultaneously occur across the swarm. These can, as shown in section 3.1.5, successfully be covered using the proposed satellite swarm approach. In case of monolithic satellites, this is obviously not valid. Design flaws are, given the low cost nature of swarm satellites more likely to occur as less time will be spent in the design phase, in order to save cost. The larger number of satellites however ensures a form of standardisation, perhaps even cross-platform, which then in turn allows gradual de-bugging of systems, with each generation slowly improving in quality. Also, the likelihood of a design flaw being noticed during pre-flight testing of swarm satellites is higher, as many more satellites are being tested in close succession, or even simultaneously. An increase in the number and gravity of design flaws is therefore not a given. Internally redundant systems are also highly timing-dependent, as synchronicity is imperative in order for a redundant system to function properly. In case of a satellite swarm, synchronicity might be less relevant, as timing- and communication delays are inherent to the system. Hence if, and only if, proper countermeasures against timing- and communication errors are made, the swarm should be robust against such issues.

As already established (Engelen, et al., 2012), internal redundancy for a satellite swarm would result in an increased element lifetime. The overall lifetime of the swarm does not truly benefit from additional internal redundancy, and the incurred cost penalty will in most cases prove prohibitive. For satellite swarms therefore, it would appear beneficial to solely rely on external redundancy, which has been proven to improve the system reliability, whilst focussing primarily on avoiding common-mode failures. The exact sizing of the swarm can then be determined using the methods and tools outlined in section 3.1.

4 DESIGN OF A SATELLITE SWARM ELEMENT

Designing a satellite is an iterative process. Even more so for a satellite swarm, as it requires mapping of a global objective onto the design of an individual swarm element, which, since these satellites interact with each other, can give rise to a phenomenon commonly referred to as “emergent behaviour” (Kornienko, et al., 2004), (Rouff, et al., 2004). This behaviour is difficult to predict, and it affects the behaviour of the swarm as a whole. For satellite swarms, unpredictable behaviour is seen as a risk. Thus, a design method which identifies the possible outcomes of such behaviour was proposed in (Engelen, et al., 2011), allowing designers to decide whether the induced risks are acceptable.

Moreover, emergent behaviour could prove beneficial in certain cases, in which case it makes sense to attempt designing the mission goals around a swarm exhibiting a given emergent behaviour, which then in turn leads to a spacecraft design which, when used in a satellite swarm, will exhibit such behaviour.

4.1 COMMON SATELLITE DESIGN METHODS

Traditionally, satellites are designed through following a top-down design process (Wertz, et al., 2011), (Kirkpatrick, et al., 2009), which is sometimes referred to as the process of decomposition (Buede, 2009). A top-down design process, strictly speaking, involves decomposing the mission requirements into requirements for the various systems and sub-systems, until finally arriving at the lowest level; after which each of the systems and sub-systems are produced and verified according to the requirements. Space system designers have grown accustomed to largely applying this method due to the unavailability of standardised off-the-shelf components. Since satellites remain a niche application, for which standardisation and mass-production of identical systems, which is commonplace for industrial and bulk-market components, has in most cases not been economically viable. Satellites are generally also one-of-a-kind and, due to the high risk of failure, are generally designed in evolutionary steps, which make a classical top-down approach suitable.

In most satellite projects different development teams have different responsibilities, satellite components are generally produced as distinct, monolithic sub-systems, which are then combined into the full satellite during a process called integration. These processes, which were at times applied too strictly, as well as the increasing availability of potent COTS components, lead to the recent boom in applying a bottom-up design approach, in which the available off-the-shelf systems are combined in a manner which suits the application. Bottom-up design approaches are gaining momentum mainly for the design of low-cost, highly miniaturised satellites (e.g. (OHB-Sweden AB, 2012)), as more and more standardised off-the-shelf components become available. This creates a form of modularity, yet adds overhead due to an excess in interfaces, cable harnesses and housings required to do so. In the case of satellites following the CubeSat-standard (Nugent, et al., 2008), (California Polytechnic State University, 2013), some of this overhead is reduced due to the application of a common (outer) housing and shared power conditioning circuitry.

Each of these methods has their merits and drawbacks, and both are rarely applied literally, as iterations are nearly always required. For the design of a satellite swarm

element, certain aspects differ from those for traditional satellites. Examples are the impact of volume-production of the units and the impact of emergent behaviour on the overall system. In case of a satellite swarm, the “space segment” can be seen as a virtual satellite, consisting of a number of physical elements, which only become effective through cooperation.

4.1.1 Top-down swarm satellite design

A traditional top-down systems engineering approach would design the swarm based on the number of elements required to achieve a certain coverage or observability. In a subsequent step, the element functionality would be designed, based on an initial assumption on payload requirements as well as the number of payloads in orbit. The resulting swarm in this iteration features the minimal number of elements required to perform the mission, and all payload data is to be aggregated in order to meet the specifications. (Temporary) loss of a single element however cannot be tolerated as no extra elements are defined, which effectively nulls some of the advantages of a swarm. A certain over-definition of the number of satellites in the swarm would clearly solve the issue, but quantifying the required number of extra satellites cannot be performed in advance, as it would require detailed system simulations using as much input data on the exact element specifications as possible, which is not available in a first iteration.

4 Like most engineering processes, top-down systems engineering therefore becomes more effective when the process is treated as an iterative process. Further iterations would then take into account the reliability and availability of the previous swarm satellite design and reiterate, finally resulting in an optimal swarm, and an optimised swarm satellite design. In case of a strict application of a top-down process, the design of the satellites will not make use of off the shelf components, as those will never exactly meet the specifications defined by the top-level (or derived) requirements. Each of the satellites will therefore consist entirely of tailor-made components, which potentially drives up the cost of the mission.

4.1.2 Bottom-up systems engineering methods

A bottom-up approach would design elements in order to meet the payload criteria, and subsequently define the number of satellites to meet the observability criteria. This however generally results in an over-definition of the swarm, as payload pooling and payload data augmentation through (mathematical) combination of different data streams, such as through applying methods as compressed sensing or image composition (Scherzer, 2011), is not taken into account. Compressed sensing techniques for example apply sparsity to reconstruct incomplete data-sets, perhaps correcting for missing parts, lost due to the failure of a single swarm satellite.

As a result, the designed swarm is also the most robust and simple, as it does not rely on element-to-element interactions to perform its basic mission. This implies that even a single swarm element could perform observations, allowing the swarm to either be built up gradually (Verhoeven, et al., 2011) or to remain operated until all of the elements failed. The satellites can be designed using as much off-the-shelf components as possible, in order to reduce the unit costs. A certain penalty would follow, in that off-the-shelf components are not tailored to the application, and will therefore always exceed the minimal required performance somewhat. Re-iterating the design using the same strict methods will not easily result in a more optimal satellite design, as most of the changes

would be driven by the availability of off-the-shelf components which meet the mission criteria, and the swarm will therefore always over-perform. A second iteration however could include processes such as data pooling or compressed sensing, as information on the performance and number of satellites present in the swarm would be available at this point. The resulting swarm however would remain slightly over-defined; yet the application of COTS components could result in a lower overall cost compared to a swarm designed through a strict, traditional top-down design process. This is mainly true for swarms designed for the minimum number of elements, in which the potential mass-savings offered by customised components does not offset the cost-advantages offered by fully off-the-shelf components.

4.1.3 Comparison

The strengths and weaknesses of both methods are shown in TABLE 4-I. This allows for a direct comparison between both methods, trading off speed of convergence of the solution, optimality and total cost and mass of the solution, as summarised in TABLE 4-II. In this table, TS(n) represents a numbered strength ('S') of the top-down method. TW(n) represents a weakness ('W') of the same method, whilst the bottom-up method's strengths and weakness use a 'B', instead of a 'T'.

These lists are not exhaustive, but indicate the key strengths and weaknesses of both methods. Cost is estimated without taking into account bulk-rate cost reductions. Even for custom components, ordering larger quantities will reduce the unit costs. Taking into account the fact that fully custom solutions are close-to-ideal, the total cost of the system designed using a top-down method could potentially be lower than that designed using a bottom-up strategy, as for example total mass of the solution could be lower for a solution using non-standardised components, due to the potential for reduction in overhead. In contrast however, COTS components benefit from being supplied to various other missions simultaneously, so the bulk rate cost reductions can be spread across different missions, and they can gain flight experience more easily. Hence no firm advantage can be defined in terms of overall cost. Both methods would have to be tried for a select system in order to verify which solution would be the most cost-effective for that particular system. Going through both design procedures in parallel however also adds a cost penalty, as two competing designs would emerge from this procedure, of which only one would be used. It would therefore seem wise to try and find whether a given method outperforms the other for a given design, eliminating the need for having a competition for each new design.

The summary in TABLE 4-II shows that the bottom-up method converges more quickly onto a workable solution. The bottom-up solution also has the additional benefit that most components are off-the-shelf, so their design-, production- and qualification times can be subtracted from the overall time required to design and produce the satellites. Hence in terms of time-to-market, a bottom up solution would be the quickest, pending availability of the components. The solution rendered by such a method is far from optimal however, yet could still prove to be more cost effective, due to the lower unit costs and development costs. This mainly holds when the number of elements remains low, as otherwise the overhead in terms of development time and cost caused by using a fully customised solution diminishes. Another drawback of the top-down method is that the risks of flying unproven hardware are higher, which generally results in additional precautionary measures being implemented into the system; which in turn increases the overall cost.

TABLE 4-I

LIST OF THE STRENGTHS AND WEAKNESSES OF STRICT TOP-DOWN AND BOTTOM-UP ENGINEERING METHODS (MAKING USE OF COTS COMPONENTS) WHEN APPLIED TO THE DESIGN OF SATELLITE SWARMS

Top-down		Bottom-Up	
STRENGTHS			
TS1	Custom designed components are always optimised to the problem	BS1	Use of off the shelf components offers a lower total mission cost
TS2	A close-to-optimal solution emerges after only a few iterations	BS2	The first iteration already ensures the satellites will meet the mission criteria, as well as the number of satellites is sufficient
TS3	The first iteration can already include effects of data-pooling	BS3	Reliability and effective performance data on the components is instantly available, through the use of off-the-shelf components
		BS4	The resulting swarm is robust and simple, as cooperation is optional because of the over-design of the elements
WEAKNESSES			
TW1	No flight heritage is available for the components, which increases the risks. Also no reliability data is available at the first iteration, and at further iterations, only estimates are available	BW1	COTS components over-perform, as they are not tuned to the application
TW2	The first iteration under-defines the number of elements	BW2	More than one iteration is only useful when pooling effects are to be taken into account, or when new/updated COTS components are available
TW3	Loss of an element is not acceptable, with the design following from iteration 1, so more iterations are required in order to increase robustness	BW3	Mass of the final solution can be higher, due to the application of COTS components
TW4	Cost of custom components is generally high; potentially resulting in a higher total mission cost		

The top-down method however will, additional risk avoidance measures aside, deliver a better solution in terms of mass and performance. The time-to-market can be higher though, as more iterations are generally required to arrive at an acceptable solution, and custom(ised) components have to be developed.

TABLE 4-II

INDICATIVE COMPARISON OF THE STRENGTHS AND WEAKNESSES OF STRICT TOP-DOWN AND BOTTOM-UP DESIGN METHODS, WHEN APPLIED TO THE DESIGN OF A SATELLITE SWARM ELEMENT

Property	Beneficial characteristics	Detrimental characteristics
Speed of convergence of the solution	BS2, BS3, BS4, BW2	TW1, TW2, TW3
Optimality of the solution	TS2, TS3	BW1, BW3
Total cost of the solution	BS1	TW1, TW4
Mass of the solution	TS1	BW3

4.2 ALTERNATIVE DESIGN METHOD

Systems engineering for satellite swarms is a new field as no actual engineering and flight experience is available. It differs from traditional systems engineering due to the fact that in a satellite swarm multiple elements cooperate to form e.g. a larger instrument. This cooperation however is relatively loosely defined as it can vary significantly for different applications, and elements will not always be able to comply with the demands of the swarm due to (internal) problems or simply because they are in the wrong physical location.

A swarm may be regarded as a monolithic larger satellite with intermittent internal communication links. Moreover, the swarm's emergent behaviour will affect its global behaviour. Emergent behaviour, which is a product of the interactions between the elements (which may be intermittent), must be predicted in order to prevent the emergence to result in undesired behaviour. Emergent behaviour becomes increasingly difficult to predict when the size of the rule-sets in the individual elements increases, complicating the design of such a swarm. A proper systems engineering method, specifically tuned to satellite swarms, is therefore required to account for these deficiencies. However, due to the variable nature of the system and the large number of variables involved, no engineering method will ever be perfect. This implies the systems engineering method *itself* will have to include contingencies for deficiencies, as not all variables can be foreseen in advance, and no simulation can be run which includes *all* contingencies. This can be minimised through extending the design process, yet time-to-market will at some point also play a significant role, at which point an extended design phase seems counter-productive.

An alternative, hybrid design approach was therefore proposed in (Engelen, et al., 2011), which attempts to combine the advantages of both methods; limiting the over-definition of the bottom-up approach, whilst quantifying the number of extra satellites at an earlier stage compared to using a top-down approach. The method relies on behavioural and observability simulations throughout the design process. This requires continuous updates throughout the process and as a result is highly iterative. It does lend itself to automatization however, which could reduce the time required for the increased number of iterations, and shares similarities to methods applied in concurrent design facilities (Bandecchi, et al., 1999), (Winner, et al., 1988), which run design processes for different aspects of a given system in parallel, based on progressing estimates of the system's parameters.

4.2.1 Element design procedure

A general overview of the method proposed in (Engelen, et al., 2011) is shown in Fig. 4.1. The same method was also shown in Fig. 3.1, which focused specifically on designing the element behaviour in order to meet the swarm's functional and behavioural requirements.

Initially, traditional mission analysis should provide the desired capabilities of the swarm as a whole. Those requirements are then fed into the design of the separate elements. Simulation and design analysis becomes an integral part of the systems engineering process, as the emergent behaviour and the global properties of the swarm are defined by more than the specifications of the individual elements alone. The output of those simulations, which apply some of the methods outlined in Chapter 0, is verified against

the global requirements, and the element design, and their behavioural rules are adapted if required.




Fig. 4.1: A global overview of the proposed satellite swarm element design method

However, the simulation is only an approximation of the actual system and even using hardware-in-the-loop simulations, accounting for all unforeseen issues will prove difficult. In order to account for the imperfections of the simulations, a certain amount of engineering margins are required to provide robustness of the system. In extreme cases, some level of human intervention might even be desirable, for example as external (independent) observers in a majority voting process.

Going more into depth into the element design process, it becomes apparent that their design process resembles that of certain low-cost nano-satellite design practices (Jackson & Epstein, 2000). For fast, relatively low complexity designs, nano-satellites can be developed using a bottom-up approach. The spacecraft's requirements are defined initially, based on best estimates of what would be achievable given the current state of technologies available on the market. During the design process, components are selected based on the most suitable off-the shelf component available. Only on rare occasions or for particular mission-critical payloads a fully custom-built solution is to be designed, as this lengthens the design process considerably. It is this approach which allows nano-satellites, and in particular highly standardised CubeSats, to be developed in a very short timeframe. Most large spacecraft are fully custom-built solutions. This shows in their high quality of output, but sadly also in their relatively lengthy development times. Standardisation in larger spacecraft is in process however, as seen for example in many standardised spacecraft busses available at various satellite suppliers. Spacecraft swarms would therefore benefit highly from applying a similar approach, yet two distinct differences should be taken into account.

Firstly, spacecraft swarm satellites are likely to be produced in larger quantities, implying that care will have to be taken to allow for a degree of mass-production of the craft. Highly integrated subsystems and modular, standardized systems are therefore highly useful for spacecraft swarms, and are therefore to be recommended. This also imposes that more extensive (per batch-) testing will have to be performed, to assure common-cause failures are not present in the design for example. Designing for mass-production also implies some form of yield management is highly desirable which could result in design-changes which improve the manufacturing speed and/or the yield.

Secondly, the much desired and/or feared emergent behaviour of a swarm is not a function of the elements themselves (Engelen, et al., 2010). More specifically – if the elements are to operate in the most efficient manner possible, they should try and use the emergent behaviour to their advantage (Kornienko, et al., 2004). It would make sense then to design a single, simple and elegant set of behavioural rules, in order to simplify analysis of the resulting swarm behaviour (Rouff, et al., 2004). However, the simplicity of their sets of rules should not allow for conflicts to occur, as emergent behaviour could be both beneficial as well as detrimental.

The design process shown in Fig. 4.1, combined with the overall behavioural design process shown in Fig. 3.1, is an attempt to reduce the impact or likelihood of just these issues. Common mode failures can be taken into account into the system-wide swarm behaviour simulations. Issues or conflicts resulting from the interactions between the elements, and hence faults in their rule sets should be discovered through these simulations as well. Note that in order to simulate communication and/or processing delays or communication faults, hardware-in-the-loop simulations are required at one of the final iterations of the design process, in order to verify any assumptions made on their likelihood of occurrence as well as the type of disturbances encountered.

4.2.2 Impact on the swarm design

The high-level swarm design method, as detailed in Chapter 0 is centred on the behaviour of both the swarm, and its elements (Engelen, et al., 2011). This is in contrast to traditional systems engineering (both bottom-up and top-down) which focuses mainly on functionalities, rather than behaviour, as without emergent behaviour, a monolithic system will behave as it was designed. In the case of satellite swarms, functionalities, while still important, are secondary to the overall behaviour of the swarm, as swarm elements will not always be able to fulfil a specific function, for example due to communication issues, elements being (physically) in the wrong location or in some fault-recovery mode, or even simply malfunctioning. This is also the case for traditional satellite systems, yet the high degree of autonomy desired when operating a satellite swarm as effectively as possible implies the swarm itself should address such issues. The element behaviour then defines, based on a pre-defined rule-set, what actions the element, and in response the swarm, takes in such cases.

The proposed design method (cf. Fig. 3.1) attempts to address the issues caused by emergent behaviour. It initially applies a high-level, top-down design method for the swarm, combined with a thorough requirements definition. The swarm's elements are designed in parallel with the swarm system, using a bottom-up approach. The initial high-level process output remains straw man-like, as the final design will only crystallise after several process iterations. The process flow of the method is shown in Fig. 4.2.

4

This method allows for a time-efficient element design, whilst simultaneously keeping track of the effects of the emergent behaviour. The hope is that this emergent behaviour can be tuned such that the output of the swarm is increased when compared to the summed output of the individual elements. Emergent behaviour is therefore a desirable effect, yet also difficult to predict and possibly chaotic, implying safeguards are useful in preventing system lock-ups or other potentially catastrophic malfunctions. Note that the initial design method for the global swarm behaviour is defined as a top-down approach, yet after the first iteration performed on the element design it morphs into a hybrid approach, adapting to the output of the system simulations, as was detailed in Fig. 4.1.

As Fig. 4.2 shows, elements and behavioural rules can be defined in parallel. However, as the communication layer is designed at a rather late stage, iterations are required to assess the impact of the available bandwidth on the global behaviour of the swarm. Also, when assuming distributed communication- and control methods are applied, for example using gossip-like communication protocols, the overhead caused by the swarm control will be largely independent of the number of elements. This is particularly beneficial for larger swarms as their control communication overhead would be significant, if it wasn't distributed or localised. Alternatively, local communication can be confined in smaller sub-clusters, with only communication between each of the sub-clusters (Budianu, et al., 2011).

This whole process can then be iterated, as shown in Fig. 4.1 and Fig. 3.1, finally resulting in a swarm of satellites that satisfies the mission requirements with a minimal number of excess elements, and therefore minimised mission cost. It is however imperative that most of these design steps are automated. Iterations are required, as the reliability of the individual elements, as well as their lifetime and individual availability, can only be determined properly when the element design is finalised for that particular design iteration. Given that those are input parameters for the swarm sizing processes, iterations are required. As with many iterative processes, initially these parameters are likely to be based on a "best guess".

Architecturally, the distinction between a swarm element and a regular satellite can remain rather small – swarm elements may differ only e.g. due to the addition of a (short-range) inter-swarm communication link. The behaviour control and the related safeguards can be located in the control software present in their on-board computers. As section 3.2 proposed, internally redundant systems are superfluous in the case of swarm satellites, as all of the swarm satellites are redundant (spare) copies of each other. This implies that another distinction would be that the internal architecture of a swarm satellite can remain comparatively simple.

4.2.3 Verification

Using the design method proposed in this Chapter, the required swarm performance properties are defined based on the mission requirements. These can then be decomposed to individual element performances. An example of such a process is discussed in section 3.1.2 which deals with predicting the lifetime and subsequent reliability figures of a given element for a given element design. These requirements can then be used to validate (in the case of a bottom-up design process or in case of the design process proposed in this Chapter) or define (in case of a top-down design process) the design of an individual swarm satellite.




Fig. 4.2: The process flow for the proposed design method
Adapted from (Engelen, et al., 2011)

Given that the effective performance of a swarm satellite will affect the overall performance of the swarm, and given the difficulties of performing in-situ repair operations, verification during the design process, and later on validation of the element's performance parameters is required. Additionally, the decomposed individual element performances will have to be verified, in order to assess the viability of a satellite swarm consisting of elements with the defined properties. This can however be assessed prior to or independently of designing the individual elements. A graphical overview of this part of the design process is shown in Fig. 4.3.

Verification of the behaviour and performance of the designed swarm and its elements, which can be defined as “proving the resulting element design complies with the overall mission requirements” can be difficult, as long-term predictions of a novel system are, at best, complicated. Accelerated life tests and other on-ground tests of the individual components form an essential part of this process. Once performances are verified on component-level, and later on element-level, system wide simulations (preferably with hardware-in-the-loop) should prove the swarm’s compliance to the mission specifications. This should be done both on a functional level, as well as on a reliability level. When considering the reliability assessments, component failure rates can be inserted directly into a Markov model for the satellite, replacing the initial ‘best guess’ (c.f. TABLE 3-1). This can then be translated into a swarm satellite reliability figure, and later on into the overall swarm reliability figures, as was demonstrated in Chapter 0.

Doing this for each of the iterations in the design process however can prove to be costly, given the amount of time and man-power required. Optimisation of the design process is therefore desirable. Certain time-consuming tests for example could be skipped in subsequent iterations or even performed only once for a specific off-the-shelf component, if those tests have proven that the sensitivity of a given design to a certain parameter change proves to be negligible. Reliability and lifetime assessments on the other hand could be integrated in an automated software-based assessment tool which computes reliabilities based on the updated inputs generated in each of the iterations. Due to the automation of this step, no significant additional delays would be posed by this assessment.

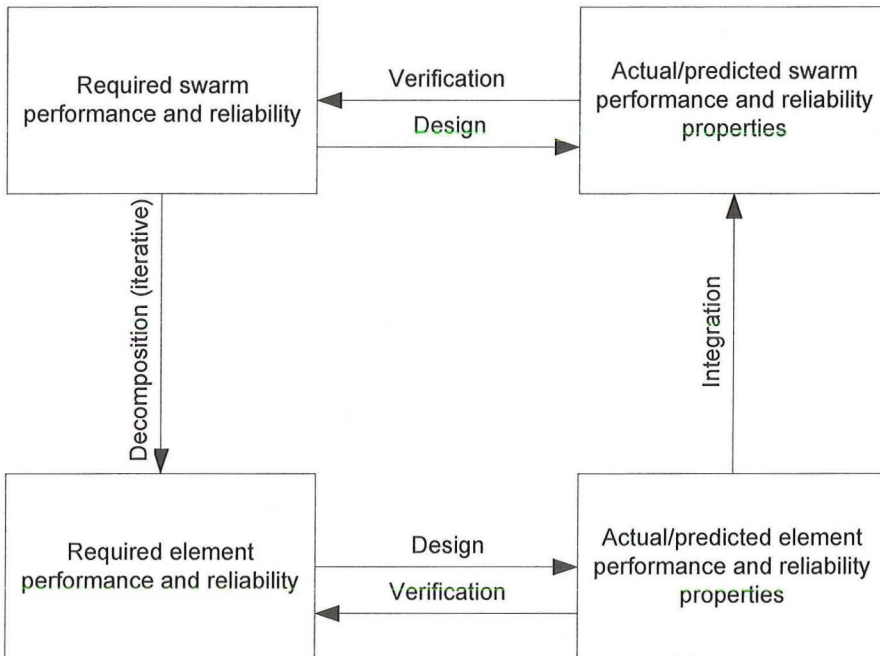


Fig. 4.3: Verification process of a given swarm design

For satellite swarms, also the behaviour should be assessed. Of primary importance are simulation and verification of an individual satellite's responses to external inputs, which in reality originate either from within the swarm (e.g. commands from other satellites), or from the environment (e.g. solar eclipses). Then, the response of the satellite will determine the actions, and by extension the reactions of the other satellites in the swarm.

5 THE OLFAR SWARM

OLFAR, short for “Orbiting Low Frequency Antennas for Radio Astronomy” (Bentum & Boonstra, 2009), aims to deploy a swarm of satellites in a remote location, in order to observe signals in a frequency range of 0.3 to 30 MHz. Observing these signals opens up the last under-observed frequency regime, allowing for example studies of the aurora’s on exo-planets and other low frequency phenomena. OLFAR plans to use an interferometric array of antennas each of which compute their correlation matrices in space. The matrices for each of the samples are then transmitted back to ground-stations on Earth; which can then distribute these to scientists for analysis.

Since Earth’s ionosphere is opaque for radio signals in these frequencies, an observation of these signals requires a space-based receiver. Also, due to ionospheric distortions, interference and man-made noise, these satellites would have to orbit at quite a large distance from Earth, in order to maximise the signal-to-noise ratio. Given sufficiently long integration-times, signals from the universe’s so-called “Dark Ages” can be received, and with a sufficiently high number of satellites, it would become possible to create sky-maps of the very early universe using such a system.

The OLFAR swarm, as described in a.o. (Engelen, et al., 2010) and (Dekens, et al., 2014) would form a schooling swarm, in a clustered spherical configuration when sampling. To date, and to the best of the author’s knowledge, this is the most advanced proposal for a schooling swarm, and it will therefore serve as the primary case study for this Chapter.

5.1 SCIENCE CASE

The field of radio astronomy started in 1931, when the radio engineer Karl Jansky discovered noise emanating from the Milky Way at a frequency of 20.5 MHz (Jansky, 1933). Jansky originally thought this noise emanated from the sun, though after prolonged observations, concluded the signals had to originate from outside of the solar system.

In 1944, Dr. Hendrik van de Hulst predicted that neutral hydrogen emits a very distinct spectral line at 1420.4058 MHz, which is referred to as the 21-cm hydrogen line (HI). This frequency can travel through optically opaque clouds of dust and gas, allowing observations of optically obscured objects. More importantly, as it is the first state of ionisation of hydrogen, almost all (warm) celestial objects emit radiation at this frequency. The first observations followed in 1951 by (Ewen & Purcell, 1951) and (Muller & Oort, 1951). The first attempts of mapping the celestial sphere and our own galaxy followed in 1958 by (Rougouor & Oort, 1960). The resulting map is shown in Fig. 5.1.




Fig. 5.1: Distributions of neutral hydrogen in the galactic plane.

The centre of the Galactic system is placed at C in the centre of the image, and the sun's position is indicated by a circle. (Rougeot & Oort, 1960)

5

Since then, ground based instruments have been able to cover most of the frequency spectrum from around 20 MHz up to about 50 GHz which is the spectral band in which the atmosphere is transparent. At lower or higher frequencies, the atmosphere becomes opaque, and at frequencies lower than about 30 MHz Earth's ionosphere severely distorts the signals. At a frequency of 408 MHz, Haslam et al. created an all-sky map, shown in Fig. 5.2 which today is still used as a reference.

Space-based telescopes, such as ESA's Planck and NASA's WMAP, were able to observe in the frequency bands where Earth's atmosphere is opaque. They completed the picture all the way into the far infrared frequency regime, as shown in Fig. 5.3 and Fig. 5.4.

At the lower frequencies however, different science cases are expected, as different physical processes are causing emissions at these frequencies (Jester & Falcke, 2009).




Fig. 5.2: The Radio Sky at 408 MHz with a logarithmic scale from 10 to 250 K (Source: (Haslam, et al, 1982))




Fig. 5.3: The cosmic microwave background at 94 GHz taken by the WMAP satellite (Source: (NASA / WMAP Science Team))




Fig. 5.4: The first All-Sky Map, generated by the Planck satellite generated from frequencies of 30 GHz to 857 GHz (Source: (ESA, 2010))

5.1.1 Science at low frequencies

The universe is theorised to have started at an event called the “Big Bang” at which point the universe started to expand, as is indicated in Fig. 5.5.

NASA's Wilkinson Microwave Anisotropy Probe (WMAP) mission mapped one of the last stages of this initial expansion, which can be observed through the microwave radiation emitted by the primordial plasma which filled the universe at that point in time, called the Cosmic Microwave Background (CMB) radiation. This map, as shown in Fig. 5.6, shows that the early universe was extremely homogenous (Bennett, et al., 2013), with thermal variations of only $\pm 200 \mu\text{K}$.




Fig. 5.5: A brief overview of the expansion of the universe
(Image credit: NASA)

After this point in time, at about 379,000 years after the Big Bang, the universe kept expanding, which cooled the plasma, which in turn causes the emitted radiation to drop in frequency. This phase lasted for between 150 and 800 million years, at which point the first stars ignited, allowing for observations in the optical or high-frequency radio domain. One of the major questions this poses is how the universe became anisotropic enough to allow condensation of plasma to form stars and galaxies. The point at which the stars ignite is called the “Epoch of re-ionisation”, and the gap in time between the images we have of the CMB and the ionisation of the first stars is referred to as the “Dark Ages of the Universe”, as no observable light was emitted at that point.

Observing the very early universe, and with it, the Dark Ages, implies observing signals with a very high relative red-shift z , which is defined as:

$$z = \frac{\lambda - \lambda_0}{\lambda_0} = \frac{f_0 - f}{f} \quad (5.1)$$

in which λ_0 is the originally emitted wavelength, and conversely, f_0 the emitted frequency, and λ and f are the received wavelengths or frequencies.

Given that the HI-line occurs at a relatively low frequency of 1420 MHz, observations of the Dark Ages ($z = 30-1000$), as well as the Epoch of Re-ionisation ($z = 6-20$), will have to take place at very low frequencies (Jester & Falcke, 2009).

Earth's ionosphere, as well as the magnetospheres of large planets, inside as well as outside of our solar system could be detected, allowing studying of the planet's aurora (Jester & Falcke, 2009). The interstellar medium, which distorts the signals in a similar manner as the ionosphere can also be studied tomographically, rendering detailed information on its structure (Jester & Falcke, 2009).




Fig. 5.6: Variations in the Cosmic Microwave Background
 The image shows a temperature range of $\pm 200 \mu\text{K}$ (red/blue)
 (Source: (NASA / WMAP Science Team))

Perhaps most importantly, this frequency regime has remained largely unexplored, hence new discoveries of yet unknown physical processes occurring in space or even on Earth are possible.

5.1.2 Low frequency radio astronomy

Low frequency radio-astronomy has more or less disappeared after the initial discoveries made in the 1960's, as higher frequencies showed less background noise, and allowed generation of much clearer sky maps due to the higher directivity achievable through parabolic dish antennas. Achieving the same at lower frequencies is hard due to the extremely long wavelengths involved, which in turn, due to practical limitations, would require extremely large antennas, which are limited to a diameter of approximately 300 m (Wilson, et al., 2009).

In 2010 however, the LOFAR antenna array (van Haarlem, et al., 2013) proved that a phased array of low frequency antennas could render the required baseline lengths for sharp, high resolution observations at low frequencies.

Observing with a synthesis array of antennas however brings its own set of problems and peculiarities which have to be addressed in order to allow generation of an accurate image. Primarily, aperture synthesis relies on obtaining a number of unique baselines, with a baseline consisting of the vector between two stations which performed two synchronised simultaneous recordings. Correlation of these two recordings then renders a single "observation". Such observations can be entered into a so-called visibility matrix, which can be said to adhere to the van Cittert – Zernike relation (Eq. (5.2)) (van Cittert, 1934), (Zernike, 1938)), due to the long distances involved

$$V(u, v) = \iint I(l, m) e^{-2\pi i(ul+vm)} dldm. \quad (5.2)$$

In this equation, u and v represent the array coordinates, which form a plane perpendicular to the vector towards the source, and l and m are the direction cosines, which are two orthogonal angles formed between the (x, y) -plane near the source, and the plane of the array. Using this relation, it becomes possible to compute the source brightness distribution in the two dimensions, given here by the intensity function

$I(l, m)$, allowing astronomers to create a 2-dimensional image of the source. Fig. 5.7 shows an overview of the angles and coordinate systems used.




Fig. 5.7: Geometry and coordinate systems used in synthesis imaging (Thompson, et al., 2001)

The relation shown in Eq. (5.2) is a simplified version of Eq. (5.3), as it assumes a 2-dimensional co-planar array, with a narrow field of view (Carozzi & Woan, 2009), through the assumption of small angles.

$$V(u, v, w) = \iint I(l, m) e^{-2\pi i(ul + vm + w(\sqrt{1-l^2-m^2}-1))} \frac{dldm}{\sqrt{1-l^2-m^2}} \quad (5.3)$$

In case of a small array, with a small field of view, l and m are small enough such that the term given in Eq. (5.4) can be neglected, compared to the $(ul + vm)$ component of the equation.

$$\left(\sqrt{1-l^2-m^2}-1\right)w \approx -0.5(l^2+m^2)w \quad (5.4)$$

Given a sufficient number of observations, a map of all sources present in the sky can be reconstructed, through a de-convolution process (Högbom, 1974). The process of reconstructing the intensity matrix and then de-convolving the resulting “dirty-beam” is commonly referred to as “all-sky imaging”. For most earth-bound antenna arrays, the

van Cittert – Zernike relation is simplified through assuming a 2-dimensional, co-planar array, with a narrow field of view. This holds for most high-frequency parabolic dish arrays. For LOFAR however, these assumptions are not always valid (Carozzi & Woan, 2009) resulting in artefacts in the resulting images. For LOFAR, most of these artefacts can be eliminated through careful post-processing procedures, as well as through the use of different imaging methods (Tasse, et al., 2012).

The net effect of removing the simplifications caused through the assumption of a co-planar array however is a significant increase in the required amount of computations to be performed on the raw data-set (Yashar & Kembal, 2009). This is a process which for LOFAR requires the use of a super-computer, and a so-called “FX correlator” (van Haarlem, et al., 2013). Data at the telescopes is gathered locally, and then either beam-formed or integrated locally. The resulting data streams are then transported to the super-computer for Fourier transforming (the “F” in “FX correlator”), correlation (“X”) and post-processing. Finally, the correlated data is deconvoluted in the imaging process.

For such an array to function, accurate timing is imperative, which is why each of the LOFAR stations features GPS-corrected rubidium atomic clocks (van Haarlem, et al., 2013). Even so, the small drift of these clocks has to be corrected for, as the drifts do show in the resulting images.

5.1.3 Low frequency radio astronomy: From space

Earth's ionosphere is a layer of plasma, which distorts electro-magnetic waves passing through it (Spoelstra, 1997). At low frequencies, the ionosphere acts as a refractive medium which (amongst others) causes shifts in the observed location of the source. This distortion shows in signal delays, which affects for example GNSS signals (Lejeune & Warnant, 2008), but also any attempts at imaging. Moreover, the ionosphere also emits radio waves at and below its plasma frequency. For Earth, the plasma frequency of the ionosphere lies around 10 MHz during the day and 5 MHz at night, which makes observations below these frequencies impossible (Jester & Falcke, 2009).

For higher frequencies, mainly frequencies above 50 GHz (Liebe, 1983), Earth's atmosphere is opaque, which implies that observations at these frequencies or above have to be done above the ionosphere. Also moisture and particles suspended in the air will trouble observations in these frequency ranges. This is the main reason for the Wilkinson Microwave Anisotropy Probe (WMAP) and Planck satellites to have been launched into the Sun-Earth L2 point which has the added advantage of being at a distance of 1.5 million kilometres from Earth, effectively nulling all man-made interference.

In 1968, NASA launched their Explorer 38 satellite which was also known as the Radio-Astronomy Explorer A (RAE-A) (Weber, et al., 1971). This satellite (shown in Fig. 5.8) was intended to perform observations of astronomical sources from an Earth-orbit, yet rather unexpectedly it discovered that Earth emitted very strong long wavelength signals, now known as the Auroral Kilometric Radiation (AKR) (Grabbe, 1981). These signals were unexpectedly intense, severely troubling any planned observations of galactic sources with RAE-A.




Fig. 5.8: NASA's Explorer 38 satellite.
(Colorado State University, 2013)

NASA therefore launched RAE-B in 1973 (Alexander, et al., 1975), into a lunar orbit, hoping that the moon would shield Earth's interference, whilst counting on the large distance between the satellite and Earth reducing Radio Frequency Interference (RFI)

signal strengths. Since RAE-2 orbits further away from Earth, it was also deemed able to sample down to lower frequencies, as in lunar orbit, the lower frequency limit is imposed by the interstellar plasma, rather than the local plasma in Earth orbits (Weber, et al., 1971).

This was the case, as Fig. 5.9 shows, and RAE-2 was able to create the first all-sky maps at low frequencies (Novaco & Brown, 1978). These maps, one of which is shown in Fig. 5.10 was made using a single antenna, using the moon as a shield, to date are still the best maps available at these frequencies.

In case an interferometer is flown in space, one must consider the fact that Earth isn't blocking signals originating from the back of an antenna anymore, and generally the w-component is not negligible or fixed anymore. This then requires solving the full visibility-equation (Eq. (5.3)), which requires significantly more processing power, as well as communication bandwidth. Various imaging techniques have been compared by (Yashar & Kembal, 2009), in light of the SKA-project; which will have to deal with significant w-components in their measurements as well.

One such example is the RadioAstron mission (Kardashev, et al., 2013) which was launched in 2011 into a highly elliptical high Earth-orbit for performing extremely long VLBI experiments with ground-based telescopes. The maximal achievable baseline between the Spektr-R spacecraft and the ground-based telescopes amounts to 330 000 km, allowing for very high resolution studies of various sources.




Fig. 5.9: The effect of occultation by the Moon on signals originating from Earth. (Alexander, et al., 1975)




Fig. 5.10: The all-sky map at 1.31 MHz, as generated by RAE-2. (Novaco & Brown, 1978)

Recently, several studies have been performed, assessing the viability of a low frequency interferometric array in space. One of the first studies performed was the DARIS study (Boonstra, et al., 2010), (Saks, et al., 2010)). The concept used in that study was to make use of space-qualified, proven hardware to form an interferometer, observing in the frequency bands between 1 and 10 MHz, with an instantaneous observing bandwidth of 1 MHz. The final outcome was that it would be possible to provide an array of nine nodes, when using a centralised FX-correlator architecture (Rajan, et al., 2013). In such a system, the antenna data is correlated in space, after which it is transmitted to Earth for post-processing, calibration and imaging.

DARIS would be placed either in an Earth-leading or -trailing solar orbit, or in a lunar orbit, in order to reduce the amount of Earth-based RFI. The DARIS nodes would fly with baselines between 15 metres and 100 kilometres, limited by the instrument's diffusion limit, caused by the presence of the inter-stellar medium (ISM). Several science cases were lined up for DARIS. They are presented in TABLE 5-1. In this table, $\delta\Omega$ represents the required angular resolution of the instrument, which then links to the required baseline for the observations, and τ represents the required integration time, at 5σ standard deviation. As can be seen, direct observations of the cosmic Dark Ages are difficult to achieve, as they require integration times in the order of years, and preferably well over 1000 antennas nodes (shown in the column titled N_{ant}), as this limits the required total integration time. In fact, many of the science cases would benefit from a higher number of nodes. Due to the high mass of the DARIS satellites, which were estimated at 100 kg per satellite node (and consequently their large volumes), this would not be economically feasible for the mission cost of 500 M€. Even so, the science cases a DARIS-like system would be able to cover appear to be more than interesting enough, shown by the large number of follow-up proposals having been written (e.g. (Baan, 2012), (Klein-Wolt, et al., 2012), (Bentum, et al., 2009)).

TABLE 5-1
 REQUIREMENTS FOR VARIOUS TYPES OF OBSERVATIONS,
 ADAPTED FROM (JESTER & FALCKE, 2009) AND (BOONSTRA, ET AL., 2010)

	Frequency [MHz]	$\delta\Omega$	N_{ant}	Baseline length [km]	τ (5σ)
Cosmology					
- Epoch of re-ionisation, Global signal	30-150	$2\pi Sr$	≥ 1	0	2h-30yr
- Epoch of Reionisation, tomography, spatial mapping	30-150	$\sim 10^\circ\text{-}1'$	$\geq 10^3$	0.7-20	\sim yr
- Dark Ages, Global signal	30-45	$2\pi Sr$	≥ 1	0	\sim yr
- Dark Ages, 21-cm power spectrum	30-45	$\sim 20^\circ\text{-}2'$	$\geq 10^4$	1-20	\sim yr
Extragalactic surveys	10	$1'$	300	0.1-100	2 yr
Galactic surveys					
- Solar system neighbourhood	0.1-10	$O(\circ)$	10-100	0.3-30	\sim yr
- Cosmic rays ¹⁶	0.1-30	$1''$	10^5	$3.10^3\text{-}30.10^3$	100d
Transients					
- Solar, planetary bursts	0.1-30	$O(\circ)$	10-100	0.5-200	min-hr
- Extra-solar bursts	0.5-30	$\leq 1'$	$\geq 10^4$	$\geq 35\text{-}1000$	min-hr

One of these is called OLFAR (Orbiting Low Frequency Antennas for Radio astronomy) (Bentum & Boonstra, 2009). In contrast to DARIS, OLFAR proposes to use as much COTS hardware as possible, in the form of a swarm of nano-satellites (Engelen, et al., 2010). The intention is to reduce the mass- and the cost per antenna-node. In return, this would allow launching many more antennas at a similar mission cost as a DARIS-like mission. OLFAR would observe in the frequency band between 0.3 and 30 MHz with an instantaneous bandwidth of 1 MHz (Engelen, et al., 2010). Since LOFAR is capable of observing down to 20 MHz (van Haarlem, et al., 2013), OLFAR would have an overlapping region, which could aid LOFAR in compensating for the significant disturbances caused by the ionosphere at these frequencies.

The use of a satellite swarm would allow for a significant expansion of the number of nodes in the array, which, given sufficient funds, could allow for tomographic observations of the Dark Ages of the universe: one of the most important observables in this frequency regime. This would however, as shown in TABLE 5-1, require over 10^4 antenna nodes, observing for many years. This then places significant constraints on the requirements posed on the OLFAR nodes. First of all, they would have to be quite reliable and highly autonomous, in order to allow trouble-free continuous observations for a long period of time. The array should also be replenished at regular intervals, as individual satellites will age and fail over time. It is therefore imperative for the OLFAR nodes to be as low cost as possible to produce, launch and operate. Perhaps most importantly, OLFAR should be located in a very quiet environment, as the signal-to-noise ratio ultimately determines the required integration time for any given observation and any reduction in integration time will result in significant cost savings.

¹⁶ Only valid for low lunar orbits

OLFAR could also serve as a pre-cursor for a lunar telescope array. Such an array would allow for continuous observations; yet at a significantly higher cost due to the complexity of the logistics of placing and operating a radio telescope array at the far side of the moon.

Observations with a space-based array, especially for an array with large numbers of antennas have never been performed to date. It requires solving the 3D-imaging equation for a non-coplanar array, which is computationally intensive. Recently, sparsity and compressed-sensing approaches are being investigated, which show the potential for a significant reduction in the amount of computations required (McEwen & Wiaux, 2011), (Wenger & Magnor, 2010). For large numbers of antennas, the inter-satellite link bandwidth becomes a practical limitation as to what is achievable for a given system (Rajan, et al., 2013). Distributed correlation solves the bandwidth requirement for the down-link (which sends the correlation matrices to the processing station(s) on Earth), but requires a vast amount of data transfers between all of the antenna nodes, as each of the individual baselines requires correlation. Systems such as OLFAR would therefore benefit tremendously from also distributing the correlation effort. This, to date, has not been achieved however.

Most, if not all, low frequency radio astronomy mission concepts intend to move to a remote location in space, in order to avoid the RFI caused by Earth's ionosphere and man-made emissions which penetrate the ionosphere. (Klein-Wolt, et al., 2012) provide a table comparing the measured RFI levels at different locations in the Earth-Moon system. This table is repeated and expanded here, using measurements performed by Cassini and WIND/WAVES in TABLE 5-II.

TABLE 5-II
RFI LEVELS MEASURED IN SPACE, AND THE REQUIRED EQUIVALENT BIT DEPTH

Mission	Location	Date	Frequency [MHz]	Max. RFI level measured [dB]	Equivalent bit depth required [bits]	Source
RAE-2	Lunar orbit	1970's	1-10	30-40	4.98 – 6.64	(Klein-Wolt, et al., 2012)
WIND/ WAVES	HEO (200000 km)	1994	6.125	45	7.48	(Kaiser, et al., 1996)
		1994	10.325	40	6.64	(Kaiser, et al., 1996)
FORTE	LEO (800 km)	1997	38	40	6.64	(Klein-Wolt, et al., 2012)
		1997	130	-10	1.00	(Klein-Wolt, et al., 2012)
Cassini	Flyby at 1186 km	1999	2-16 MHz	25	4.15	(Fischer & Rucker, 2006)

As can be seen, the RFI levels drop significantly with increasing frequency. The RAE-2 measurements were taken for a lunar orbit, not taking the immersion into the radio-quiet zone into account. Attenuation levels of 10-30 dB are shown to be achievable in low frequency ranges (Jester & Falcke, 2009), (Alexander, et al., 1975) in the 1334 × 1123 km lunar orbit in which RAE-2 orbited the moon. As shown by (Takahashi, 2003), a lower lunar orbit is beneficial as the attenuation will become higher. The moon itself also imposes a lower frequency-limit of approximately 200 Hz to the observations, limited by its size.

5.2 OLFAR SWARM IMPLEMENTATION

The scientific mission of OLFAR requires a large number of antenna nodes to fly in a remote location in space, with a maximal baseline of 100 km, in a spherical configuration as defined in (Engelen, et al., 2010). This baseline requirement is confined by the presence of inter-planetary medium, which permeates the solar system. Surveys at a frequency of 10 MHz show that the worst-case angular scattering caused by the inter-planetary medium amounts to approximately one arc-minute, limiting the maximal baseline length to approximately 100 km (Dekens, et al., 2014), (Jester & Falcke, 2009).

The OLFAR telescope is also aimed at being confusion limited¹⁷; which at 10 MHz amounts to a sensitivity of 65 milli-Jansky (Jester & Falcke, 2009). This sensitivity requires an observational time of approximately 5.2×10^4 days (Dekens, et al., 2014), when observing with a single OLFAR-like satellite node. Increasing the number of simultaneous observations, through increasing the number of satellite in the swarm, linearly decreases the required observational time. Should the OLFAR telescope consist of 50 satellites for example, the total observation time required to reach the confusion limit will reduce from 142 years for a single satellite to 2.9 years for 50 satellites.

The relative timing between these nodes is to be as close to ideal as possible, as otherwise the images lose coherence already during their required integration time (Rajan, et al., 2013). Their relative ranges have to be known to within $1/10^{\text{th}}$ of a wavelength, in order to allow determining the length and orientation of the baseline. In this particular case, the shortest wavelength amounts to about 10 metres for 30 MHz. This results in a requirement on the relative position knowledge of one metre. The process of clock-synchronisation can be performed using the inter-satellite link, which simultaneously allows for determining the relative range between all of the nodes within the network (Rajan & van der Veen, 2013). Absolute knowledge of the position of the array is irrelevant as concerns the science observations.

The required range-rate is limited by the observational integration time. For OLFAR, it is assumed each of the individual satellite nodes integrates each measurement for one second, in order to reduce the total amount of data to be transmitted. Given that most astronomical sources are continuum sources this is not considered to be an issue. The snapshot integration time of one second however limits the allowable amount of baseline drift during the measurements, as the assumption is that all of the samples are taken in a quasi-static scenario. This limit is set at $1/10^{\text{th}} \lambda$, which in turn limits the allowable relative velocity of each of the satellites to 1 m/s.

The OLFAR antenna nodes are, ideally, swarm satellites, as a satellite swarm's self-management property would allow for an extremely large number of satellites operating in close proximity to each other. Given that the orbits of the satellites do not demand very high control accuracies (contrary to the determination accuracies which are higher, with a requirement of relative ranging to within one metre accuracy), allowing the swarm to manage itself appears plausible.

The scientific requirements for the OLFAR telescope, presented initially in (Bentum, et al., 2009), and later updated in (Rajan, et al., 2011) are collected here in TABLE 5-III.

¹⁷ The confusion limit is the sensitivity at which celestial objects cannot be distinguished from unresolved objects.

TABLE 5-III

PRIMARY REQUIREMENTS FOR OLFAR, DERIVED FROM (BENTUM, ET AL., 2009) AND (RAJAN, ET AL., 2011)

Frequency range	0.3-30 MHz
Antennas	Dipole or tripole
Number of elements	>10, scalable
Maximum baseline between satellites	100 km
Spectral resolution	1 kHz
Processing bandwidth	100 kHz
Spatial resolution at 1 MHz	0.35 degrees
Snapshot integration time	1 s to 1000 s, depending on deployment location
Sensitivity	Confusion limited
Instantaneous bandwidth	1 MHz
Deployment location	Earth orbit, Moon orbit, Earth-Moon L2, Earth-Moon or Sun-Earth L4/5

5.2.1 Orbit design

One particular orbit scenario was initially deemed the most suitable for OLFAR due to its relative proximity to Earth. This scenario was the lunar orbit case, and it has been extensively studied, in close cooperation with and under supervision of the author, for its suitability as a science orbit for OLFAR (Dekens, et al., 2014). Other equally viable alternatives are an Earth-Moon L2 orbit, or perhaps even an Earth-Moon L4 and L5 orbit. Orbits considered for DARIS included Earth-leading or Earth-trailing solar orbits as well, which for the case of OLFAR are considered as problematic due to the severe distance to be bridged by the communication systems. High Earth-orbits are also still possible, in case the RFI signals are sufficiently predictable. Such an orbit would require a larger inter-satellite link capacity, as the sampling bit-depth will have to increase to allow mitigating the effects of RFI. The close proximity to Earth however increases the available throughput to the ground stations, which could increase the overall scientific data products delivered in the course of the mission.

Low Frequency astronomy in lunar orbit

The Moon is shown to have interesting properties for use in low frequency astronomy. (Alexander, et al., 1975) have shown that the Moon effectively shields most of the RFI originating from Earth or Earth's ionosphere, as shown in Fig. 5.9. The Moon itself will also interact with extremely highly energetic particles. These interactions are likely to generate radio-bursts which could be studied by the system. This would effectively convert the moon into a large particle detector (Jester & Falcke, 2009).

The eclipse fraction can be approximated by a simple exercise in geometry, as the moon shields signals originating from Earth, as shown in Fig. 5.11. In this figure, r_E represents the radius of Earth, r_{Ea} represents the radius of Earth's atmosphere and r_{EI} represents the radius of Earth's ionosphere. r_M represents the radius of the Moon, r_{SWO} the radius of the swarm's orbit, and α is the angle between the Earth-Moon orbital plane and the edge of the eclipse cone, given by $\alpha = \tan^{-1} r_M / d_{apex}$. The mean distance between the Earth's and the Moon's centres is represented by d_{EM} . The eclipsed area is shown shaded in grey.




Fig. 5.11: Geometries involved in determining the radio eclipse fractions for lunar orbits

The diameter of the eclipse cone at the height of the orbit of the swarm can be calculated using Eq. (5.5):

$$R_{cone} = \frac{R_M(d_{apex} - r_{swarm\ orbit})}{d_{apex}} \quad (5.5)$$

in which

$$d_{apex} = \frac{d_{EM}r_M}{(r_S - r_M)}$$

In this equation, $r_{swarm\ orbit}$ represents the radius of the orbit from the lunar centre, and r_S the radius of relevant noise source with respect to Earth's centre. TABLE 5-IV then tabulates the apex distances for a few relevant noise sources.

TABLE 5-IV
APEX DISTANCES FOR VARIOUS NOISE SOURCES

	Symbol	Radius of noise source [km]	Apex distance [km]
Earth's surface	r_E	6,370	144,311
Ionospheric and atmospheric noise	r_{Ea} or r_{Ei}	6,970	127,755
LEO satellites	-	7,870	108,998
GEO satellites	-	42,156	16,532

As expected, signals originating from the ionosphere have the largest effect on the size of the cone. This effect increases as the orbit height of the swarm increases.

The fraction of the orbital period spent in eclipse can be calculated as:

$$F = \frac{\sin \alpha (d_{apex} - h_{orbit})}{\pi(r_M + h_{orbit})} \quad (5.6)$$

This is tabulated in TABLE 5-V, for a few candidate orbital heights, at 0° inclination between the Earth-Moon orbital plane.

TABLE 5-V
ECLIPSE FRACTIONS FOR SELECT CANDIDATE ORBITAL HEIGHTS, FOR LUNAR ORBITS

Noise source:	Surface	Ionosphere	LEO	GEO
Orbit height [km]	Eclipse fraction			
200	28.51%	28.50%	28.49%	28.05%
500	24.64%	24.63%	24.61%	23.85%
1000	20.07%	20.05%	20.03%	18.89%
2000	14.60%	14.58%	14.54%	12.95%
3000	11.44%	11.41%	11.36%	9.51%
5000	7.93%	7.90%	7.84%	5.70%

As can be seen, ionospheric noise will eclipse less than man-made noise sources on the surface. Geostationary satellites, should they generate noise in these frequency bands, are the most influential, as the eclipse fraction defines the maximum total observation time available per orbit¹⁸. As expected, increasing the orbital altitude reduces the eclipse fraction. Increasing the orbital altitude also increases the orbital period, implying that the lowest lunar orbit achievable should be targeted as this would maximise the total observation time. Lower lunar orbits however place stringent demands on the thermal management of the satellites. As will be shown later, orbital maintenance for lower orbits increases as well, yet to an extent this is beneficial to the mission.

Whilst this radio-occultation is proven to exist, its exact properties are still unknown (Jester & Falcke, 2009), due to, amongst others, the presence of the moon's (tenuous) ionosphere (Klein-Wolt, et al., 2012). This causes a frequency-dependent refraction, which is thought to diffract signals with lower frequencies more than signals with higher frequencies. The net effect for a lunar mission is that the radio-silent cone is not as sharply defined, with low frequencies still present at the edge of the cone, as shown in Fig. 5.12 for two select frequencies. Certain lunar-surface based missions aim at using the Malapert mountain to shield their instrument from Earth-based RFI (Klein-Wolt, et al., 2012).

¹⁸ Provided no observations are performed when not in radio eclipse. Observations outside of the eclipse cone are possible; yet RFI signals will be worse.

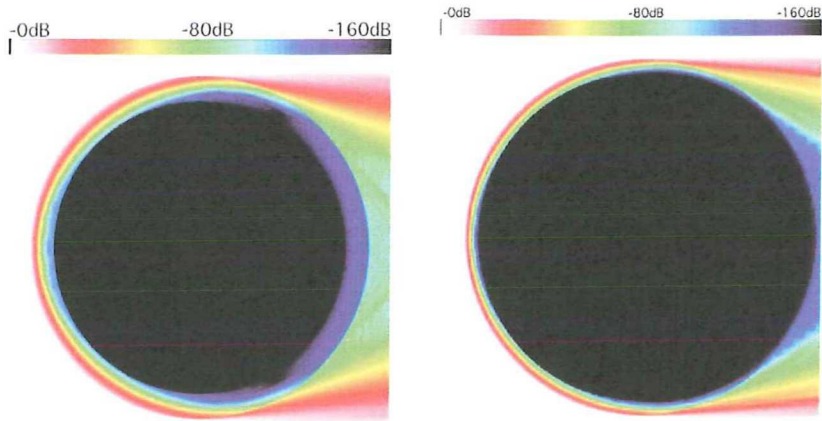


Fig. 5.12: Energy density distribution around the Moon with a continuous 30 kHz (left) and 60 kHz (right) plane wave incident from the left. (Takahashi, 2003)

Furthermore, there are also non-technical issues which complicate science with a lunar orbiting array. For one, an agreement between many radio-astronomers prevents communication at the far-side of the moon, as it is intended to be designated as a radio-quiet zone (Maccone, 2008). This in turn severely increases the demands on the inter-satellite link capacity, as well as the coherence time of the satellite's internal clocks, as no communication, and hence no synchronisation is allowed for a significant fraction of the orbital period.

Reference science orbit

Lunar orbits are notoriously instable, due to the large variations in mass concentrations on the Moon. This poses problems for lunar orbiters, as they require orbit maintenance manoeuvres in order to compensate for these effects. In the case of OLFAR however, the exact orbit is not considered relevant, and variations in the orbital altitude can be accepted. This was studied by Erwin Dekens, under the supervision of the author, and published in (Dekens, et al., 2014). It was shown that the orbit evolves in a cyclical manner, as shown in Fig. 5.13.

As can be seen in the figure, the apo- and periselene vary by over 80 km over a period of about 200 days. For lunar-observing missions this would be dramatic, yet for OLFAR, the only criteria would be to remain in orbit (i.e. never drop below a given altitude limit to avoid collisions with surface features). The perturbations caused by the lunar gravity-field also have advantages, in that they aid the process of filling the UVW-sphere. This means that satellites can scan different points in the sphere without applying orbit-corrective manoeuvres.




Fig. 5.13: Orbital evolution for an initial orbit of 200 km propagated over a period of five years. (Dekens, et al., 2014)

In (Dekens, et al., 2014), it is shown this period lies at approximately 100 days for a 200 km orbital altitude; at which point the swarm drifts outside the 100km inter-satellite distance boundary. After this period, the swarm would have to apply a single corrective manoeuvre, reverting to dense cluster. They also investigated a higher orbit at 3000 km, which shows an evolution as shown in Fig. 5.14.




Fig. 5.14: Orbital evolution for an initial orbit of 3000 km propagated over a period of five years. (Dekens, et al., 2014)

As can be seen, the orbital variations caused by the lunar gravity-field are much smaller. In fact, the dominant disturbance force experienced in this orbit is the third-body attraction by Earth. The coherence-time, defined as the time it takes for the swarm satellites to drift too far apart, is also much higher for this orbit. The downside is that scanning is less effective; which means more satellites are required to provide the same amount of UVW-coverage in an equal amount of time.

Swarming

A reference orbit, inclined at 5.148° with respect to the ecliptic, which equals 0° inclination with respect to the Earth-Moon orbital plane can be selected based on the eclipse fraction offered, and the total amount of science time available per year. This orbit can then be used as a reference, around which the other satellites can form a spherical swarm, with relative distances of up to 50 km in each direction, effectively forming baselines of up to 100 km. Such orbital patterns are achievable for a lunar orbit, as was shown in (Dekens, et al., 2014), through varying the relative orbit parameters of each of the swarm satellite's orbits. Comparisons were made between an optimised orbital distribution of the swarm, versus a random distribution. An example of the UVW coverage after a single orbit of the optimised configuration is shown in Fig. 5.15.




Fig. 5.15: Baseline paths in UVW-space covered by a 25-satellite swarm after one orbit. (Dekens, et al., 2014)

As can be seen, a significant number of baselines are already covered, even in a single orbit. The eclipse fraction has not been included in this graph however, limiting the useful baselines. However, random distributions were also compared, as shown in TABLE 5-VI.

TABLE 5-VI
SIMULATION RESULTS FOR 36 DISTINCT SWARM SCENARIOS, FOR A PERIOD OF 100 DAYS.
SOURCE: (DEKENS, ET AL., 2014)

(a) Randomized configurations

Scenario				Results	
#	Initial in-plane baseline limit	Orbit altitude	Number of satellites	F_c	F_d
1	0 km	200 km	100	0.771	0.351
2			25	0.293	0.351
3			5	0.031	0.351
4		3000 km	100	0.002	0.123
5			25	0.001	0.123
6			5	0.000	0.123
7	50 km	200 km	100	0.824	0.350
8			25	0.345	0.350
9			5	0.040	0.350
10		3000 km	100	0.220	0.123
11			25	0.062	0.123
12			5	0.003	0.123
13	100 km	200 km	100	0.933	0.347
14			25	0.487	0.347
15			5	0.056	0.348
16		3000 km	100	0.741	0.121
17			25	0.138	0.121
18			5	0.006	0.121

(b) Optimized configurations

Scenario				Results	
#	Initial in-plane baseline limit	Orbit altitude	Number of satellites	F_c	F_d
1	0 km	200 km	100	0.991	0.307
2			25	0.823	0.250
3			5	0.112	0.248
4		3000 km	100	0.405	0.123
5			25	0.213	0.123
6			5	0.014	0.123
7	50 km	200 km	100	0.995	0.232
8			25	0.900	0.183
9			5	0.104	0.164
10		3000 km	100	0.924	0.117
11			25	0.403	0.109
12			5	0.022	0.109
13	100 km	200 km	100	0.981	0.333
14			25	0.747	0.279
15			5	0.088	0.326
16		3000 km	100	0.351	0.123
17			25	0.182	0.122
18			5	0.009	0.123

In this table, F_c denotes the UVW-coverage, and F_d denotes the mean duty cycle, imposed by the radio eclipse fraction. The duty cycles reported in TABLE 5-VI are higher than the approximate values shown in TABLE 5-V, as they are more exact. Diffractive and refractive effects are also not taken into account however. The table was generated in order to allow comparison of 36 distinct scenarios for optimising the swarm distribution around the reference orbit. Parameters such as the relative anomalies or relative eccentricities were varied. More details on this process can be found in (Dekens E. , 2012). The results for a 5-satellite swarm were averaged over 20 simulation runs, whilst the results for the 25-satellite swarms were averaged over four runs. Due to compute-power limitations, the 100-satellite swarm results were not averaged.

As the analysis shows, randomised distributions are outperformed by the optimised configuration in all cases. The effect however diminishes with an increasing number of satellites in the swarm, which seems to validate the proposition to use a satellite swarm given the large numbers of antennas desired for OLFAR, rather than a formation flight of satellites. It is also apparent that the 3000 km reference orbit is less suited to natural UVW-scanning, resulting in lower coverage figures.

In (Dekens, et al., 2014), a full simulation of scenario #13 was made, which varied the relative anomalies of the satellites, for an in-plane baseline limit set at 100 km and an orbital altitude of 200 km, using 100 satellites. As TABLE 5-VI shows, this scenario results in a coverage of the UVW-space of 93.3%, with a mean duty cycle of 34.7%, due to the differences in the orbits of the satellites. Fig. 5.16 shows the result of the simulation, over the course of the entire 100 day period, for a selection of cross-sections of the UVW-sphere. In this scenario, the voxels were assumed to have a size of 1 x 1 x 1 km, in order to limit the amount of memory consumed during the simulation runs. As can be seen, the central volume of the UVW-sphere is covered very frequently, implying that

short baselines are easy to come by. The edges are quite sparsely covered after 100 days, indicating that the longer baselines are quite rare. Certain volumes are not covered at all in this period which shows that even 100 satellites are not sufficiently covering the sphere, for a random initial distribution. Should the swarm, after this period of 100 days, contract in a completely different “initial” configuration, it is not unlikely that these voxels would also be covered. The total integration time at these baseline lengths however remains extremely limited; which is one of the primary reasons for requiring extremely large numbers of satellites for tomographic studies of the Dark Ages, as already shown in TABLE 5-I.

In (Dekens, et al., 2014) the potential of this scenario in terms of achievable coverage of the UVW-space, the cumulative measurement time as well as the highest relative baseline-rate was also analysed. It was found that for the 200 km reference orbit achieving the 52,000 days of cumulative observation time required to reach the 65 mJy sensitivity can be achieved after 365 days with 25 or more satellites in the swarm, or using 100 satellites in the 3000 km reference orbit.

After a period of 100 days, 49% of the UVW-sphere is covered using 25 satellites in the 200 km orbit, and 93% using a swarm of 100 satellites. For the 3000 km reference orbit, 74% can be achieved using a swarm of 100 satellites, after a period of 100 days. With less satellites, it appears to be imperative that the satellites apply active corrective manoeuvres after 100 days; not only to realign themselves to remain in the 100 km sphere, but also to inject some additional offsets in their orbits, which would allow scanning other voxels in the sphere. The findings are summarised in TABLE 5-VII.

TABLE 5-VII
ACHIEVABLE PERFORMANCE FIGURES FOR THE TWO CONSIDERED LUNAR REFERENCE ORBITS.
(ADAPTED FROM (DEKENS, ET AL., 2014))

	Achievable value						Requirement
	200 km reference orbit [nr. of satellites]			3000 km reference orbit [nr. of satellites]			
	5	25	100	5	25	100	
Baseline coverage [%]	> 6 ^C	> 49 ^C	> 93 ^C	> 1 ^C	> 14 ^C	> 74 ^C	≥ 95 % ^A
Cum. Measurement time [days]	2600 ^B	77000 ^B	1300000 ^B	880 ^B	26000 ^B	440000 ^B	≥ 52000 ^A
Cum. Measurement time [days]	712 ^C	21095 ^C	356164 ^C	241 ^C	7123 ^C	120547 ^C	≥ 52000 ^A
Highest baseline rate[m/s]	116	116	116	30	30	30	≤ 3

^A After the full mission duration

^B Extrapolated to one year of mission duration

^C After 100 days of free drift

The main issue, identified in (Dekens, et al., 2014) is the relative velocity of the satellites. Satellites with the largest baseline of 100 km will experience a relative velocity of 116 m/s for a 200 km lunar orbit, or 30 m/s for a 3000 km reference orbit. This is linked firmly to the orbital mechanics, and can be calculated using Eq. (5.7), assuming Keplerian motion, in which B represents the overall desired maximal baseline distance between any two satellites, μ_{moon} the Moon’s standard gravitational coefficient (equal to 4902.8 km³/s²), and a the semi-major axis (Dekens E. , 2012)

$$\dot{b}_{max} = B \sqrt{\frac{2\mu_{moon}}{a_r^3}}. \quad (5.7)$$

This equation is valid for orbits which use spreading of the relative anomalies of the orbits for creating the swarm. Orbits using relative eccentricities to create the required baselines show less relative velocity (Dekens, et al., 2014)

$$\dot{b}_{max} = B \sqrt{\frac{5\mu_{moon}}{4a_r^3}}. \quad (5.8)$$

As shown in (Dekens E. , 2012), varying the relative eccentricities could prove to be worthwhile, as it offers similar degrees of coverage, while resulting in slightly reduced relative velocities. The initial limit for the baseline velocities originates from the requirement that in order to perform correlation, the physical location at which the sample is taken is not allowed to move more than, ideally $\lambda/10$, in order to guarantee full phase coherence. At worst, each sample is to be acquired within $\lambda/3$, which equates to a baseline rate of 1 m/s or at worst 3 m/s, when assuming an integration time of 1 second, at the highest observation frequency of 30 MHz. Alternative schemes which apply different baseline lengths for different frequencies can be imagined; though the scaling with baseline length is linear. Hence halving the baseline, which coincidentally also halves the achievable instrument resolution at that wavelength, only results in a 50% reduction of the baseline rate.

Tabulating the maximal baseline rates (see TABLE 5-VIII) versus orbital height shows that such rates are only achievable for reference orbital heights of over 15000 km, which do not show a sufficient degree of natural disturbances to allow efficient sampling with a randomised swarm. While this would be possible using a satellite swarm, this would plead for using a formation flight of satellites in this case. The duty cycle, and hence the cumulative observation time at such altitudes, as well as the degree of attenuation offered by the Moon at these altitudes is rather limited however, with duty cycles in the order of 2.9% at 15000 km. Alternatively, the observation frequency could be reduced to 1 MHz for example, which would allow for a baseline rate of 100 m/s worst case.




Fig. 5.16: Filling of the UVW-voxel sphere for scenario #13.

The colour of the voxel indicates the number of times the voxel is covered by a baseline-pair of satellites, in which blue indicates a low number of passes, and red a high number.

(Dekens E. , 2012)

TABLE 5-VIII
 BASELINE RATES IN LUNAR ORBIT FOR VARIOUS ORBITAL ALTITUDES,
 USING TWO DISTINCT METHODS FOR GENERATING SWARM ORBITS

Method:	Relative anomaly variations	Relative eccentricity variations
Altitude [km]	Baseline rate [m/s]	Baseline rate [m/s]
200	116.1	91.8
500	93.6	74.0
1000	69.2	54.7
2000	43.3	34.3
3000	30.4	24.0
5000	17.9	14.2
7500	11.2	8.8
10000	7.8	6.2
15000	4.6	3.6
20000	3.1	2.4
25000	2.3	1.8

Performance of a lunar science orbit

As shown, a lunar orbit for an OLFAR-like swarm shows great promise. The lunar gravity field induces natural random drifts in the satellite orbits, which benefits filling of the UVW-sphere. A satellite swarm without any strictly defined orbital distribution would quite easily meet the science objectives, saving on the number control manoeuvres as well as propellant. As with any swarm mission, increasing the number of satellites in the swarm greatly improves the performance of the instrument, as the number of baselines scales with the square of the number of nodes in the array. Furthermore, the Moon acts, especially at lower orbital altitudes and for higher frequencies, as an effective RFI-shield.

The only, yet very significant issue with a lunar orbit are the relative velocities of the satellites, which enforce a reduction in the total integration time per sample. Alternative options to solving this issue exist, such as increasing the orbital altitude, or limiting the maximum observation frequency or maximum baseline length. These options are not desirable from a user-perspective however, nor from an inter-satellite bandwidth point of view.

The Moon itself can also double as a particle detector, which renders access to different science cases, compared to a free-flying swarm.

Other orbits

Other orbits considered for OLFAR are three distinct lunar-Lagrange points, which are points in space which appear stationary with respect to two orbiting bodies, like for example the Moon itself orbiting Earth, or very high Earth orbits. Of these, the Earth-Moon Lagrange points; namely LL2, LL4 and LL5 show the most promise in terms of RFI levels and instrument performance. An overview of the candidate orbits is shown in Fig. 5.17.




Fig. 5.17: Candidate orbits for the OLFAR array.

LL(n) represents an Earth-Moon Lagrange point, whilst MO represents a lunar orbit scenario. HEO represents a High Earth Orbit case. Note that the sizes are not to scale.

LL4 and LL5 are nearly identical copies of each other. They are stable points leading or trailing the Moon in its orbit around Earth, and are therefore not shielded by the Moon. The same holds for a high Earth orbit, which lacks the gravitational stability of a Lagrange point. High Earth orbits (HEO) are also closer to Earth, which is beneficial for the downlink capacity, yet detrimental for the sampling depth required to capture the dynamic range posed by the strong RFI signals and the weak signal of interest. Swarming around an Earth orbit follows the same strategies as swarming around a Lunar orbit however, yet lacks the gravitational disturbances caused by mass concentrations in the Moon. Given a sufficiently high altitude however, the Moon will influence the orbit, as it will act as a strong source of third-body perturbations. The relative speeds in these orbits should not pose any issues; which in turn could benefit the scientific observations, as longer integration times reduce the required bandwidth for correlation and downlink of the sampled data; which would allow increasing the sampled bit depth without introducing an extra burden on the communication links. For an HEO at an altitude of 275,000 km for example, the relative baseline rate, calculated using Eq. (5.8), would amount to 0.6 m/s, which meets the original $\lambda/10$ criterion. Given the relaxed requirement of $\lambda/3$, or 3 m/s, and an identical link budget, sampling could occur at a bit depth of 5 bits per sample, which could satisfy the RFI dynamic range at this altitude.

As the Moon isn't used as an RFI shield, sampling can happen continually, potentially increasing the total data volume gathered per year.

LL2 is a saddle-point, which allows for wide halo-orbits, near-vertical orbits or complex Lissajous orbits. The fact it is a saddle-point implies directional stability; which could possibly be used to introduce a scanning motion in the satellite orbits. The point itself lies some 60000 km behind the Moon, as seen from Earth. The Moon therefore acts a

shield, yet the net effect is difficult to predict, due to the frequency-dependent refraction which occurs near the lunar surface.

TABLE 5-IX lists the eclipse cone size at the LL2 point. As can be seen, signals originating from geostationary orbits will always propagate to any satellite in LL2, as the LL2-point lies behind the apex distance for such signals, causing the cone diameter calculation to result in a negative diameter. For these purposes however, it implies there is no shielding.

TABLE 5-IX
ECLIPSE CONE DIAMETER AT THE LL2-POINT

	Radius of interest	Apex distance	Cone diameter at swarm distance	
Earth surface	6,370	144,311.0	2,000.6	[km]
Atmosphere	6,970	127,755.3	1,808.9	[km]
LEO satellites	7,870	108,998.4	1,521.4	[km]
GEO satellites	42,156	16,532.1	-9,433.5	[km]

It should also be apparent that the cone is relatively small, implying that large halo-orbits are not shielded. Communication with Earth is troubled due to the increased distance, but mainly due to the fact that a swarm orbiting inside the cone does have a line-of-sight with the ground stations, and that communication on or near the far-side of the Moon is frowned upon (Maccone, 2008). Relay-stations placed in geostationary orbits will always have a line of sight with the satellites, provided they are positioned in the antipodal arc with respect to the position of the Moon. As an alternative, relay stations could be placed in LL4 or LL5. The main drawback for any relay station placed in space is the required antenna size: Earth-based ground-stations are not really limited in terms of aperture, giving them significantly more gain. For a space-based relay systems, this is much more expensive to achieve.

For any of these candidate orbits, the Moon itself cannot easily be used as a navigational reference point either, complicating navigation (Beliën, et al., 2011).

High Earth Orbits show an RFI level of around 40 dB (see TABLE 5-II) which would require sampling with a bit depth of around 7 bits. Even LEO orbits show similar noise figures, which would imply that an array orbiting closer to Earth would be suitable, at the expense of an increase in the amount of data to process and transfer by a factor of 7-8, compared to the original OLFAR concept requirements. The reduction in distance to Earth would allow for a faster downlink-speed however, partially compensating for the increase in sampling depth.

Relative speeds in HEO's are also acceptable, as shown in TABLE 5-X, which allow for longer integration times. In a 275,000 km HEO for example, the relative speed amounts to 0.6 m/s (worst case), which would already allow integration of the data by a factor of 5. This results in a reduction in transferable data by a factor of 5 as well, strengthening the case for a HEO science orbit. As with the Lagrange point orbits however, autonomous navigation is hindered by the lack of a nearby absolute reference point. Currently, GNSS receivers are being developed which would allow operation above the GPS constellation altitudes, possibly up to Lunar Lagrange points distances (Winternitz, et al., 2009), (Carpenter, et al., 2004). Availability of such receivers would solve both the navigation issue, as well as the clock drift problems encountered.

TABLE 5-X
 BASELINE RATES IN HIGH EARTH ORBITS FOR VARIOUS ORBITAL ALTITUDES,
 USING TWO DISTINCT METHODS FOR GENERATING SWARM ORBITS

Method:	Relative anomaly variations	Relative eccentricity variations
Altitude [km]	Baseline rate [m/s]	Baseline rate [m/s]
35768	10.32	8.16
50000	6.67	5.27
100000	2.57	2.03
150000	1.44	1.14
200000	0.95	0.75
250000	0.69	0.54
275000	0.6	0.5

5.2.2 Swarm distribution methods

A single OLFAR satellite in lunar orbit essentially has a few primary tasks. It has to collect science data, synchronise clocks and relative distances between the other satellites. Then it has to correlate its captured data to the data captured by the other satellites in the swarm, in order to form the visibility-matrix. This could be done on ground, yet the bandwidth requirement for the down-link proves to be prohibitive. That matrix should then be transmitted to a ground-station on Earth for further processing. Local RFI mitigation could also be in the task list, provided enough information on the RFI is known a-priori to allow in-situ removal.




Fig. 5.18: Operating phases for a given OLFAR satellite
 in a counter-clockwise lunar orbit.
 (Engelen, et al., 2010)

The phases shown in Fig. 5.18 are defined by the location of the radio eclipse. In this figure, activities such as recharging are considered to be autonomic, as they do not involve any form of intelligent decision making, and are therefore shown on the lowest shell (closest to the moon). Maintenance operations, such as repositioning of the element, as well as powering down non-essential systems prior to entering the observation phase, are shown on the shell above that, as they are conscious actions taken by the element, at element level. Processing and science observation activities are considered to be activities performed by the payload, in contrast to the autonomous swarm-satellite platform, and are therefore shown on a separate shell, as they are assumed to take place on a “payload level”. Each “synch” phase, during which the satellites synchronise their time, location and current payload data-set, involves inter-satellite communication, and it is therefore shown on the one-but highest shell, as it concerns an activity at “swarm level”. The downlink-process, ergo sending the processed data to Earth, is shown on the outer shell as this process takes place at a global space-segment level.

This scenario could represent the nominal operations phase of the OLFAR array, when placed in a lunar orbit. As shown in Section 5.2.1, this phase can remain stable for up to 100 days in a low lunar orbit scenario. After this period the swarm will have to reconfigure itself, in order to retain useful baseline lengths for observations. Also during this period, collisions could require the use of avoidance-manoevres, which require active orbit corrections of at least the involved elements. As also shown in (Dekens, et al., 2014), given a sufficiently high number of swarm satellites the coverage offered would be more than sufficient to meet the scientific objectives. For swarms consisting of 25 satellites for example, in a low lunar orbit, the 65 mJy target will already be reached during the first year of operations. Faster accumulation requires adding more satellites to the swarm, yet would be advantageous as orbit corrective manoeuvres require pausing the observations due to the reconfiguration of the array.

Given the close relative orbits followed by the OLFAR swarm elements, actively avoiding collisions appears to be imperative for the survival of the swarm, as a collision would create a significant amount of debris orbiting at similar velocities in similar orbits as the rest of the swarm. Collision avoidance involves predicting the position of each of the elements well in advance, and quite likely also interacting with the other members to optimise global propellant consumption. Such control mechanisms have been studied for Earth-based swarm-systems, as well as for space based satellite swarms. The topic of most of these research activities is aimed at using local controllers only, drawing inspiration from natural swarms. Global communication is also prohibited, such that the system can be scaled to large numbers of satellites, without posing additional strain on the inter-satellite links. A prominent example is a study performed at ESA (Pincirolì, et al., 2008), which analysed an equilibrium-shaping behaviour-based approach proposed earlier by (Izzo & Pettazzi, 2007) for swarms with up to 500 members.

An alternative method, based on sequential convex programming is proposed in (Morgan, et al., 2013), which focuses on optimizing propellant consumption during swarm reconfiguration manoeuvres, while minimising the error between the intended terminal states. They also show that communication is only required between nearest-neighbours, reducing the overhead posed by such a method.

Given the reliability of swarm elements it is not unlikely however that defunct satellites would also orbit inside or near the swarm, which would require some form of conscious

decision to either avoid the defective node, or to relocate the entire swarm in order to avoid a collision.

Effective and elegant swarm control remains a complex matter; especially when communication latencies (Liang, et al., 2013) and limited communication bandwidths are involved. This remains one of the open issues to be solved.

5.2.3 The OLFAR swarm

An OLFAR swarm which can meet the scientific objectives given in TABLE 5-III can be sized using the data obtained for a lunar orbit, as given in section 5.2.1. As shown, a swarm consisting of 25 satellites or more would meet the scientific objective of a survey sensitivity of 65 mJy after a period of 100 days, if the swarm orbits the moon at an initial altitude of 200 km. The UVW-sphere coverage then amounts to 49% which could be improved during the next 100 days, or by adding more satellites to the swarm. The coverage for a 100-satellite swarm amounts to 93% for example in the same time frame. Alternatively, orbit corrections could be used to improve the UVW-coverage during this 100-day period.

In terms of reliability however, the swarm should manage to remain operational for the duration of the science observations. A trade-off can be made, comparing for example a swarm of 100 satellites with a predicted minimal lifetime of 100 days, versus a swarm consisting of 25 satellites or less, with a predicted lifetime of 250 days or more, depending on the coverage, as swarm lifetime and swarm element numbers could be interchangeable. In (Dekens, et al., 2014) it was also reported that for a 200km lunar orbiting swarm consisting of merely five elements, the coverage amounts to 6% after a period of 100 days, whilst the cumulative measurement time would amount to 712 days, after 100 days of operations. Such a swarm would require an operational lifetime of 20 years in order to meet the 65 mJy sensitivity requirement. Incidentally, such an operational lifetime amounts to a 73-fold increase in duration, which amounts to a final coverage amount of 438% in case no additional scanning manoeuvres are implemented. Given the type of observations performed by an OLFAR swarm however, short-lived yet highly numerous swarm elements would improve the scientific output significantly, as the number of inter-satellite baselines increases exponentially.

Given that at this point no firm data is available on the structure and reliability of an OLFAR element, we assume reliability estimates as calculated in section 3.1.2 as an initial estimate. The 90% reliability-lifetime of a centralised swarm satellite model amounted to 7.68 years, whilst the decentralised satellite model amounted to 0.32 years, due to the increased number of states present in that model. When taking the MTTF into account, these lifetimes are increased by a factor of 10 to 70.68 years and 3.2 years respectively, yet relying on those figures implies relying on the redundancy offered by the swarm. This has been shown to be acceptable in select cases (ref. Section 0); and could therefore be used. The reliability of the element gathering the data remains lower though, so certain data-points will be lost in the process. The swarm can then be modelled as a k-out-of-m system, for each of the estimated lifetimes. This is modelled for various expected total system lifetimes for a minimal required number of operational satellites equal to $m=4$, using the models from section 3.1.2 and the result is shown in TABLE 5-XI. Note that in this case, the model used a standard exponential failure distribution. As the table shows, a significant increase in the expected system lifetime is achievable for a given initial swarm number. Intermediate replenishment is of course still possible, yet not taken into

account in this analysis. The initial lifetimes of 73 and 3 years represent the MTTFF computed for the centralised and decentralised swarm satellite model respectively, whilst the 90% reliability points of those two measurements are shown in the next two columns. As can be seen, a swarm consisting of 100 initial elements gains about 39% in terms of operational lifetime compared to a 5-satellite swarm, which becomes inoperable after a loss of the first excess element.

TABLE 5-XI
 MTTFF OF A SWARM CONSISTING OF (N-M) ACTIVE ELEMENTS,
 BASED ON CENTRALISED AND DECENTRALISED SWARM SATELLITES.

m=4	MTTFF		90% Reliability	
	73	3	7.68	116.8
n	[years]	[years]	[years]	[days]
5	73.00	3.00	7.68	116.8
10	85.17	3.50	8.96	136.3
15	91.80	3.77	9.66	146.9
20	96.37	3.96	10.14	154.2
25	99.84	4.10	10.50	159.7
50	101.43	4.17	10.67	162.3
100	102.19	4.20	10.75	163.5

According to the scientific requirements, a 100-satellite swarm would require an operational lifetime of 100 days, which can be met with any of the swarm satellite designs, even at their 90% reliability points. In fact, when starting out with 100 satellites, 75% of the satellites can be lost, before ending up at a 25-satellite swarm, which can already on its own meet the science requirements, provided it remains operational for another 250 days. For smaller swarms, the scientific requirements require a longer operational lifetime. A 25-satellite swarm in a 200 km lunar orbit for example requires 250 days of observations in order to fill the UVW-sphere due to natural drift, whilst providing 52500 days of cumulative observation time. The same swarm in a 3000 km lunar orbit would require 788 days to meet the 65 mJy sensitivity requirement, and at least 800 days in order to fill the UVW-sphere in a natural manner¹⁹. Smaller swarms require dramatically longer integration times to meet the 65 mJy sensitivity requirement, and should therefore be avoided.

Based on these initial reliability estimates, as well as the scientific performance estimates, the initial OLFAR swarm size can be defined for various cases, when neglecting the baseline rates. These are summarised in TABLE 5-XII, and will be used in the subsequent element design procedures.

TABLE 5-XII
 INITIAL SWARM SIZING FOR A LUNAR ORBITING OLFAR ARRAY,
 NEGLECTING BASELINE RATES

Orbit scenario	Minimal swarm size	Minimal mission duration	Swarm satellite model based on 90% reliability point in MTTFF
A 200 km lunar	25	250 days	Centralised
B 200 km lunar	100	100 days	Decentralised or centralised
C 3000 km lunar	25	800 days	Centralised

¹⁹ In this orbit, natural scanning due to orbit perturbations is limited. More frequent corrective manoeuvres would increase the UVW-sphere fill-rate dramatically.

5.3 OLFAR ELEMENT DESIGN

Following the preliminary sizing of the OLFAR swarm, a thorough design of an OLFAR element can be made; assuming the element design guarantees an operational lifetime of 250 days, and allows for operation within a swarm of at least 25 elements, as was defined in section 5.2.3. Given that such a design process is quite elaborate, and highly iterative only those aspects which up to now are studied in some depth will be covered in this Chapter.

5.3.1 Overview

An OLFAR element in essence is nothing more than a communicative self-supporting orbiting set of antennas. This is reflected in the component breakdown, shown in Fig. 5.19. This figure has been subdivided into 5 meta-functionalities an OLFAR element requires, namely the payload function, the element's communication and intelligence, its AOCS, navigation and locomotion and its power system. An independent, autonomous (hypothetical) end-of-life-device has been added. Such a device could remain as simple as an emergency-beacon which starts transmitting its location at an emergency-frequency; whilst powering down the transmitters of the main satellite, in order to prevent the satellite from interfering with the operations of the remaining swarm. More elaborate versions could de-orbit the element for example, yet no such devices exist to date. It is quite clear however that swarms orbiting in close proximity of each other will require a device of this type, in order to prevent damage to the operation of the swarm. In the case of OLFAR however, the relative speeds are limited by design, reducing the impact of a collision somewhat.

Each of the meta-functionalities coincides with a level of "intelligence"; similar to the shells displayed in Fig. 5.18. The power system of an element for example is considered to remain largely autonomic. Housekeeping data, such as battery status information, can still be collected, which allows for energy-balancing across the swarm. Yet unless active actions are taken by the other systems in the element, the power system will keep tracking the sun and recharge the batteries. As will be shown in section 5.3.5, and 5.3.3, the solar panels will likely double as communication antennas, which require rotating in the direction of Earth for communication purposes. The current design of the power system also only allows for coarse sun-acquisition; relying on the attitude determination sensors to provide for a more precise estimate of the current position of the sun. The autonomy of the power system is therefore not absolute, as overriding commands can be provided.

One level up from the power system are the AOCS, locomotion and navigation functions. They consist of an array of sensors and actuators, which allow the element to navigate and (re-)position itself according to the requirements defined by the collective swarm intelligence. 'DS GPS', which an abbreviation for 'Deep Space' GPS, is marked in a dotted line, as it remains uncertain whether a useful navigation signal can be detected this far away from the navigation satellites, due to the significant free-space losses and the increase in dilution of precision.




Fig. 5.19: OLFAR Element component breakdown grouped in meta-functionalities

5

The highest level of intelligence inside the element, and out of the element takes place in the ‘communication and intelligence’ functions, which contain the inter-satellite and long-range communication transceivers, as well as the swarm intelligence processors. Also an atomic clock is required in order to remain in sync for as long as possible. This clock is shown in a lighter shade, as deep space GPS could potentially replace it; in case it becomes available. The functionality of an atomic clock however remains required. The deployment functionality is also shown in a lighter shade, as this is considered to be a one-off system; only used during commissioning of the satellite, and can from that point on be shut down.

The payload itself is considered to be largely autonomous. Even the correlation effort, which is routed through the inter-satellite link is assumed to take place inside the payload section; together with data storage and data capturing. This in order to assure a clear distinction between the swarm platform, and the payload it is carrying. In the actual system, these distinctions will likely be less visible however. One such example is that during the science observations, it is likely the solar panels will be shut down in order to limit the interference generated by the power supply conditioning circuitry. In fact, during observations, most of the satellite is likely to be powered down, and the payload will simply store its data for processing at a later stage. Even the AOCS system is not required, as the orientation and position before and after entering the observation mode can be used to compute a sufficiently accurate orientation and position estimate. The swarm element’s functionalities can also be shown in a functional breakdown diagram (FBD), as shown in Fig. 5.20.

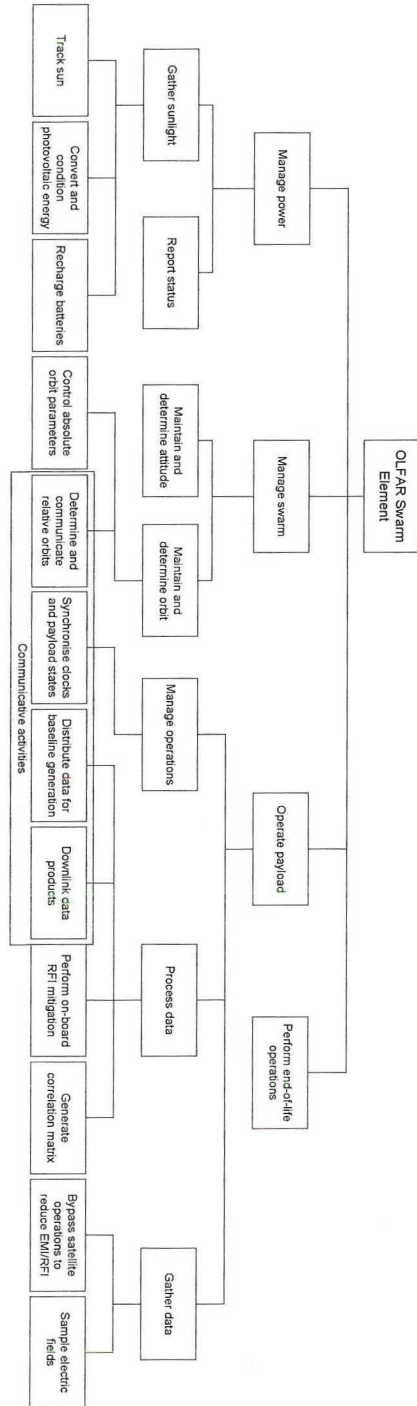


Fig. 5.20: FBD for an OLFAR element during the science phase
 As the FBD shows, four main categories of functionalities can be distinguished, specifically for the element. A fifth one, the communicative activities, is not necessarily a

discrete function of the satellite itself, as it is partly controlled by the communication originating from within (or even outside of) the swarm. This is indicated in the grey area, and it overlaps between swarm management activities, in which the element is to take part, and payload activities. It is anticipated that the payload will dominate over the swarm management activities in terms of bandwidth requirements. The satellite's functions during the science observations phase can then be sequenced into a functional flow diagram (FFD), as shown in Fig. 5.21.




Fig. 5.21: FFD for an OLFAR element during the science phase

In this figure, the three distinct parallel flows can be distinguished, which match the global functions shown in the FBD, namely a “power management”-flow, a “swarm management” flow, which also manages the element itself, and a “payload operations” flow. These processes are mostly uncorrelated, as they run independently. Certain interactions between the flows are present, and indicated in the diagram. For example, long-distance communications are assumed to use the surface area available at the back of the solar array (see Section 5.3.5). This implies that the sun-tracking mechanism will have to be disabled to allow pointing the array towards Earth. A similar dependence holds for powering down the MPPTs, and hence solar energy harvesting, during science observations inside the radio eclipse. Very little information is required flowing from one process flow to another, except for a local clock, and status information on the satellite element. This information is used to determine which of the satellites is able to assume the role of downlink-satellite, as it requires a substantial amount of battery power.

Communication items are shown in two distinct shades of grey, in which light grey implies inter-satellite communication items, and dark grey implies long-range communication activities.

5.3.2 Payload

OLFAR requires a wide-band receiver, receiving signals in the band between 300 kHz to 30 MHz. In order to reduce the correlation complexity, it is highly desirable to have three orthogonal antennas located at each satellite, as otherwise the relative orientation of the platforms has to be taken into account in the correlation process. Three orthogonal sets will allow creation of the full Stokes parameters of the incoming wave front locally, reducing the required inter-satellite bandwidth significantly.

Given the extremely wide operational bandwidth, the port impedance will vary significantly. An active antenna, which effectively uses the antenna as an electric field probe is therefore used, as it is less sensitive to the port impedance. In order to remove the side-lobes which appear in the antenna pattern at resonance frequencies, the antennas are sized to a total length of 9.6 meters, which corresponds to a resonance frequency of 31 MHz (Quillien, et al., 2013). The antennas are placed as an orthogonal set of two dipoles and two monopoles. The two monopoles are digitally combined into a pseudo-antipodal dipole due to practical constraints. Having three orthogonal dipole pairs would have been ideal, as it reduces the internal data rate by a factor of two for each of the antennas, and hence also reduces the amount of computations per sample-point. Physically this was not possible however, as the deployment mechanism of the two orthogonal dipoles was stored at that location. An overview of the geometric configuration is shown in Fig. 5.22.




Fig. 5.22: Overview of the OLFAR satellite geometry.

Antennas A and D form a dipole pair, as well as E and F.

Antennas B and A however are two monopoles. (Quillien, et al., 2013)

As can be seen in the figure, the two monopoles are antipodal, which will generate an offset in the resulting antenna pattern. This antenna pattern was studied theoretically as shown in (Quillien, et al., 2013), and later on verified experimentally as reported in (Quillien K. A., 2013) with a scale model, scaled to a frequency of 420 MHz. The tests proved the small size of the ground plane and the antipodal placement of the two monopoles would not pose significant problems.

Given the significant length of 4.8 meter per deployed antenna, a design had to be found which would minimise the storage volume of these antennas, as they are seen as driving the volume, and hence mass and therefore cost of the satellites. This has been studied extensively by (Quillien K. A., 2013), which resulted in a very capable design. A picture of the prototype of the Science Antenna System (SAS) is shown in Fig. 5.23, with antennas extruded from PPO/PS (a blend of Polyphenylene Oxide and Polystyrene) plastic. The system applies a so-called Triangular Retractable And Collapsible boom

(TRAC) to support a thin-wire antenna. The wire-antenna is sized at precisely twice the skin depth of the worst-case operating frequency. The boom is extruded in Styrene Maleic Anhydride (SMA) plastic, which proved, out of all plastics tested the most suitable in terms of manufacturability and elasticity. UV-tolerance and atomic-oxygen compatibility is still to be verified (Quillien K. A., 2013). Given that the antenna is to act as an electric field probe, it would also be possible to coat the TRAC boom with a metallic conductive layer (e.g. gold), instead of embedding the copper wire-antenna.



Fig. 5.23: Image of the SAS system prototype

5

The antennas are internally connected as two electric dipoles and two electric monopoles. Each of the antennas is connected to a low noise amplifier, after which one of two scenarios occurs: either the amplified signal is sampled directly; or an analog band-pass filter is applied prior to the sampling process. This has significant impact on the processing path, as shown in Fig. 5.24 and Fig. 5.25. In scenario A (Fig. 5.24), all forms of filtering will have to take place in the digital domain. This is therefore the most flexible solution, yet requires processing a larger amount of data, as prior to band-selection, the full bandwidth is sampled and will therefore have to be processed, placing more demands on the Analog to Digital Converters (ADC's) and band-pass filtering (BPF) hardware. It is possible however to process not one single band, but "N" bands, as the output of the band-pass filter is flexible. In scenario B (Fig. 5.25), which uses an analog band-pass filter, much less data will have to be processed in the digital domain, reducing the demands on the ADC and processing elements, at the expense of flexibility.




Fig. 5.24: Node-level signal acquisition path, with digital band-selection (scenario A)




Fig. 5.25: Node-level signal acquisition path, with analog band-selection (scenario B)

It may also prove possible to add a scenario to these two (scenario C), in which the whole band, sampled by a system identical to the one from scenario A is processed in its entirety; after which ground-based band selection can be performed. The processing chain for this is very similar to the one shown for scenario A, except that the band-pass filter is now probably not needed, and can be replaced by a low pass filter to filter out unwanted signals. The required data rates for each of the scenarios are tabulated in TABLE 5-XIII, for swarms consisting of 25 active satellites.

TABLE 5-XIII
PROCESSING PARAMETERS FOR AN OLFAR ELEMENT

	Scenario A	Scenario B	Scenario C	Unit
Number of satellites	25	25	25	-
Integration time	1	1	1	s
Observation time [†]	2683	2683	2683	s
Observation time [‡]	3221	3221	3221	s
Available processing time ^{†,C}	3450	3450	3450	s
Available processing time ^{‡,C}	20206	20206	20206	s
Number of ADC channels	4	4	4	-
Sampling depth	16	16	16	bit
Sampling rate	60000	2000	60000	ksps
Instantaneous bandwidth	1000	1000	30000	kHz
Processing bandwidth	40	40	1200	kHz
nr. of spectral channels	25	25	25	-
Number of polarisations	3	3	3	-
Number of I/Q channels	6	6	6	-
Bit depth after RFI-mitigation	1	1	6	bit
Number of frequency bins	1024	1024	1024	-
Raw sampled bitrate	3840	128	3840	Mbit/s
Raw processing bitrate	96	96	2880	Mbit/s
Required raw buffer [†]	1200	40	1200	GiByte
Required raw buffer [‡]	1440	48	1440	GiByte
Required pre-processing buffer [†]	30	30	900	GiByte
Required pre-processing buffer [‡]	36	36	1080	GiByte
Intra-satellite uplink volume, per node ^{†,B}	45	45	8097	GiByte
Intra-satellite uplink volume, per node ^{‡,B}	54	54	9720	GiByte
Total storage volume required per node ^{†,D}	1125	1125	202432	GiByte
Total storage volume required per node ^{‡,D}	1350	1350	243001	GiByte
Minimum inter-satellite link rate [†]	112	112	20160	Mbit/s
Minimum inter-satellite link rate [‡]	23	23	4132	Mbit/s
Downlink to ground station ^E	461	461	2765	kbit/s
Total volume per orbit [†]	147	147	884	MiByte
Total volume per orbit [‡]	177	177	1062	MiByte

[†] For a 200 km lunar orbit scenario

[‡] For a 3000 km lunar orbit scenario

^A Assuming a distributed correlator where $n_{channels} = n_{satellites}$

^B Assuming local integration

^C Assumption

^D Worst case, as pre-selection of correlation partners is possible

^E Per second of observation

The table is constructed assuming 16-bit ADC's sample the raw incoming signal for all four antenna channels. Each of the channels is band-pass filtered, either in the analog domain, or immediately after A/D conversion. Each raw data-stream is also integrated as soon as possible, as it significantly reduced the requirements on the local buffer. The two monopole channels are then combined into a single complex data stream, resulting in a total of three streams of complex (I/Q-) data per node. The band-pass filtering is assumed to allow distribution of as many channels as there are active satellites in the swarm, reducing the amount of data to be processed by each of the individual nodes. In total, the swarm will have processed the entire instantaneous bandwidth after recombining all individual datasets.

According to (Rajan, et al., 2013), the amount of data transferable through the inter-satellite link ($D_{intersatellite}^{fdc}$) for each of the satellites can be computed according to

$$D_{intersatellite}^{fdc} = 2N_{pol}(N_{sb} - 1)\Delta f_i \quad (5.9)$$

in which N_{pol} represents the number of polarisations to process, N_{sb} represents the number of sub-bands to process and Δf_i the instantaneous bandwidth.

The amount of processed, correlated data to be transmitted to a ground station (D_{out}^{fdc}) can then be determined using

$$D_{out}^{fdc} = \frac{2N_{sig}^2 N_{bins} N_{bits}}{N_{sat} \tau_{int}} \quad (5.10)$$

with $N_{sig} = N_{satellites} N_{pol}$, which defines the number of signals to be processed. $N_{satellites}$ defines the number of active satellites taking part in the observation, N_{bits} represents the bit depth remaining after RFI mitigation and N_{bins} represents the number of bins used in the FFT transform. Lastly, τ_{int} is the integration time used for the observation.

While immediate integration per antenna has the advantage of reducing the data rate by an amount equal to the integration time; it does remove certain rapid phenomena. It would be possible to allow a limited local round-robin buffer to maintain the raw signal for a given amount of time. These raw signals can then be scanned (automatically) for a set of pre-defined interesting features. In case a signal of interest is present in the buffer, the buffer can be stored in long-term storage devices, for transmission to a ground station on Earth later on.

It should be noted that scenario C, which aims at processing all of the bandwidth, will by far be the most challenging. Especially the inter-satellite link rate will prove problematic. The downlink rate for all scenarios is reduced significantly however, due to the limit imposed on the instantaneous bandwidth. Still, RFI mitigation is currently assumed to be perfect, which implies that only 1 bit information is sufficient for representing the measured signals. This may prove much more difficult in reality, which will increase all of the data rates shown in the table.

Other orbit scenarios are likely to relax a few of the constraints posed by the lunar orbit. Especially the long sampling-time in lunar orbits poses a problem, both on the requirements placed on long-term clock deviation, as well as on the buffer sizes required. In a HEO or LL2 scenario for example, processing and sampling could be broken up into short bursts of e.g. one minute each. The availability of longer integration times in HEO and LL2-orbits will also significantly reduce the processing and inter-satellite link loads which is why these orbits merit further investigation.

5.3.3 Communication

As shown in section 5.3.2, the payload of the OLFAR swarm results in copious amounts of inter-satellite communication, as all satellites are required to transmit their data to one another, as it is required for forming and processing observational baselines. Swarm

management communication, while required, is expected to be dwarfed by the amount of communication required by the payload processing processes.

The inter-satellite is also limited by the amount of energy reserved for it; and requires communication between the shortest baselines, as well as the longest baselines. An intelligent communication scheme has been designed (Budianu, et al., 2011), which divides the swarm in local clusters, each with a dedicated “master”, or “cluster-head” satellite. The cluster-head in this case is formed by a satellite which meets certain requirements on its available energy stored in its batteries, as well as a physical proximity to the centre of the local cluster of “slave”-nodes. The advantage is that while the cluster-head requires long-range high volume communication with the other cluster-heads, the slave nodes can save on energy. It is also not unlikely that the cluster-heads will perform the brunt of the processing, as they already have stored all of the samples taken by their slave devices. When the energy level of the current cluster-head diminishes, a slave device can then take over, as it has been able to save on energy. According to this scheme, the worst case distance a slave device would have to cover amounts to 40 km, whilst a cluster head would have a worst-case distance to bridge of 90 km (Budianu, et al., 2012). This is deemed achievable with a slave output power of 0.1 W at 2.45 GHz, whilst cluster head nodes require 4 W of output power. The achievable data rates are then 8 Mbit/s and 63 Mbit/s respectively, whilst applying 5 dBi of antenna gain; which is achievable with traditional patch antennas. A diversity-scheme has also been devised, in which the six patch antennas on all sides of the satellite body are combined to form a phased array, which increases the antenna gain in off-axis directions of the antennas, fulfilling the requirement of 5 dBi in all transmission directions (Budianu, et al., 2013).




Fig. 5.26: Render of an OLFAR element highlighting the patch antenna array

The downlink antennas (shown in Fig. 5.26) share the surface area of the solar panels (Budianu, et al., 2014), which significantly increases the link margin. Each of the solar panel segments includes four patch antennas, which are beam-formed using a binomial scheme, which maximises the gain of the array. The solar panel rotation mechanism is then used to point the antennas in the direction of the ground station on Earth. Two such panels per satellite would render a 17 dBi antenna gain. The system is predicted to

consume 25 W of power, whilst providing a downlink data rate of 900 kHz for a receiving antenna with a gain of 70 dBi.

It remains possible to modify the phases of each of the panel segments, which in turn allows for beam-forming in the direction not controlled by the solar panel rotation mechanism (the \hat{y} -direction as indicated in Fig. 5.26), eliminating the need to rotate the satellite using the attitude control system. A diversity scheme could also be applied, which is shown to provide an antenna gain of 3-5 dBi (Budianu, et al., 2014).

5.3.4 Ranging and clock synchronisation

Since all data gathered by the OLFAR array will have to be correlated both in space and time, accurate clocks are required. Outside of the radio eclipse, data-communication is allowed. A clock synchronisation scheme which also extracts relative-position information has been devised (see e.g. (Chepuri, et al., 2013)). This method has been shown to asymptotically approach the Cramér-Rao lower bound, whilst remaining extremely efficient. The method has since been expanded to include range and range-rate estimation as discussed in (Rajan & van der Veen, 2013) and (Rajan, et al., 2013).

OLFAR is assumed to cease communications whilst sampling. This implies that for the lunar orbit case, no communication occurs for up to 3300 s, in case of the 3000 km lunar orbit, which in turn requires the internal clocks of each of the satellites to remain coherent for at least this amount of time. Alternatively, in case the clock deviation can be accurately predicted, it will have to remain deterministic for at least this amount of time.

This has been studied in (Rajan, et al., 2013). They find that Rubidium-atomic clocks remain coherent with sufficient accuracy for periods up to 1200 seconds for the current chip-scale atomic clocks. Physically larger clocks can remain coherent for up to 10,000 seconds which would suffice for any of the lunar orbit scenarios. The GPS 1 Pulse-Per-Second (PPS) output is linked to the atomic clocks of the GPS system; which offers a much improved long-term stability, remaining coherent for over 100,000 seconds. The GPS signal is however not available when in Earth eclipse, which rules out this option.

To date, this problem has not been addressed by available technologies, which could result in a reduction in the sampling time. Increasing the platform size per element is another option; yet this will significantly drive up the cost of the mission, as space-grade atomic clock candidates have masses of 600 grams to 3.3 kilograms (Rajan, et al., 2013), whilst consuming between 14 and 30 Watts of power, or it will lead to a reduction in the number of elements, which in turn reduces the number of baselines.

5.3.5 Energy supply

An OLFAR element in lunar orbit experiences solar eclipses. It also experiences varying power usages depending on the operating state. Preliminary sizing of these loads has been done by (Klein, 2014), and is repeated in TABLE 5-XIV. The time estimates listed in the table are worst case times, limited only by the available time in the orbit. For the science processing phase for example, it is likely that the processing time in the 200 km orbit is insufficient to process all gathered data, whilst the time reported for the 3000 km could prove to be excessive.

TABLE 5-XIV
PRELIMINARY ENERGY BUDGET FOR AN OLFAR ELEMENT (KLEIN, 2014)

Mode	Power usage [W]	Period [hr]	Total energy consumed [Whr]
200 km lunar orbit			
- Science recording	7.6	0.75	5.7
- Solar eclipse	10	0.74	7.4
- Science processing	13.8	0.63	8.7
3000 km lunar orbit			
- Science recording	7.6	0.98	7.5
- Solar eclipse	10	1.2	12
- Science processing	13.8	5.9	81

Given that it is best to disable the MPPT's during science observations in order to reduce the self-generated RFI, batteries are required. Also, since the surface area offered by the deployable solar panels is shared by a phased array of patch antennas for use by the downlink transceiver as discussed in (Klein, et al., 2013) and (Budianu, et al., 2014) batteries are required to power the satellite when solar harvesting is disabled. A bus topology, repeated in Fig. 5.27, is proposed which is optimised for overall efficiency, based on buck-boost MPPT's. This bus topology allows for powering the satellite either directly from the solar panels in case the batteries are fully charged or during charging, or powering the satellite from the batteries in case of eclipses or when the MPPT's are disabled. The total conversion efficiency is estimated at >80%. It is furthermore assumed that the batteries are cycled once per orbit, coinciding with the science operations in a lunar orbit.

The Energy Supply System (ESS) as proposed in (Klein, et al., 2013) applies a photovoltaic array of commercial silicon solar cells, glued onto an FR-4 epoxy-based Printed Circuit Boards (PCB) as substrate, which allows integrating micro-strip patch antennas into the back-side of the substrate. These solar cells were chosen primarily for cost reasons, as the cost for the solar cells per satellite amounts to a mere 22 euro for this type of solar cells (Klein, 2014). A similar panel using space-grade cells would cost approximately 12,700 euro, though at a reduced mass, as the total panel area can potentially be reduced by a factor of 2, primarily due to the increased efficiency, and the reduced temperature sensitivity. Given the large number of satellites which are desired for OLFAR, reduced unit costs were seen as desirable, partly leading to this choice. For a low lunar orbit (a 200 km lunar orbit), these silicon cells are not suited, due to the increased temperature the cells would experience. We therefore propose a space-grade solar cell array for such a mission scenario.

The hinges of the panel are formed by Micro Miniature Coaxial (MMCX) RF-connectors, which allow feeding each segment of the panel with a unique RF-signal allowing for phase matching or beam-forming. The entire ESS is designed without redundancy. Each of the solar panel's strings is managed by its own MPPT, and each battery is managed by its own Battery Management System (BMS), which allows for graceful degradation in case of a drop in performance in any of the strings or batteries. The robustness however is assumed to be derived from the swarm, as very little excess capacity is allowed for by the design. Each panel is designed to deliver 14 W of power at the start of the science observations phase, which renders a total power per satellite of 28 W in a 3000 km lunar orbit (Klein, 2014). The Beginning of Life (BOL) power delivered by the solar arrays is much larger, at 48 W per satellite. Degradation due to traversing the Van Allen radiation belts is taken into account there. Should the thruster accept higher

power levels however, it is not unlikely that the larger BOL power allows for a faster transfer through the Van Allen belts, which in turn results in more power available at the start of the scientific operations phase.

A prototype of a solar panel was built (see Fig. 5.28) to verify the operations and test the effects of RFI emitted by the panel itself. One important finding (Klein, 2014) was that the substrate mass is significant, mainly due to the high glass content in the PCB substrate. Significant savings would be achievable through using a different and thinner substrate material, yet the RF properties will also differ, affecting the performance and design of the patch antenna array. Manufacturing issues with attaching the solar cells were also discovered. The actual level of RFI emissions emitted by the panels has not been measured yet however. Should these tests conclude that the RFI emissions are at an acceptable level with the current panel design, which includes an internal Faraday cage in an attempt to minimise these emissions, the energy budget would change entirely, as then the solar panels can remain active during scientific observations; reducing the requirements on the battery capacity.




Fig. 5.27: ESS Bus topology as proposed in (Klein, et al., 2013)




Fig. 5.28: Prototype of an OLFAR solar panel

The ESS overall, is designed for an MTTF of 15 years, and a battery MTTF of 5 years. No in-flight reliability data for this system is available at this time however; yet all known degradation factors are estimated and taken into account in this design life. It is therefore likely these failure rates will be met.

5.3.6 OLFAR element reliability

Given the current overall design of an OLFAR element, it becomes possible to assess its reliability. TABLE 5-XV lists the components, and their associated lifetime estimates. Most of this data is still taken from (Monas, et al., 2012), yet for the ESS more specific estimates are available. The atomic clock also did not have a reliability estimate, however it is assumed to be limited mainly by the read-out and control electronics. The scientific antennas, once deployed, are assumed to have an infinite lifetime. The same holds for the structure, as those failure rates are considered to be irrelevant.

TABLE 5-XV
ESTIMATED MEAN TIME TO FAILURES OF COMPONENTS
IN THE OLFAR ELEMENT DESIGN

	MTTF [yr]
ADCS	455
Atomic clock	2712
Batteries	5
Downlink	814
ESS	15
Inter-satellite link	814
Payload receiver and processor	2712
Propulsion	455
Science antennas	infinite
Storage lifetime	26.6
Structure	infinite
Swarm control processor	2712

The storage device is assumed to be similar to a commercially available Solid State Disk (SSD), which fail after approximately 5,000-10,000 write cycles for low cost devices using Multi Level Cell (MLC) type flash memory devices, and after approximately 100,000 write cycles for certain high-end Single Level Cell (SLC) type flash based devices (Thatcher, et al., 2009). The failure rate can be determined using Eq. (5.11). Note that this assumes linear write-behaviour, which is different from normal use for these devices, yet in the case of OLFAR, in which the device is used as a buffer, this is quite a realistic assumption.

$$MTTF = \frac{C_{listed}(1 + R_{spare})}{\dot{d}_{write}n_{cycles}} \quad (5.11)$$

with C_{listed} equal to the device capacity, R_{spare} the ratio of spare capacity to total capacity present in the device, and \dot{d}_{write} the write rate to the device. The number of allowed write-cycles for the device is then finally represented by n_{cycles} .

When considering these devices stop functioning after losing their excess capacity, the MTTFF can be calculated according to

$$MTTF = \frac{C_{listed}(R_{spare})}{\dot{d}_{write} n_{cycles}} \quad (5.12)$$

Given that an OLFAR element would process data at a rate of 96 Mbit/s (see TABLE 5-XIII), for approximately 35% of an orbital period (when considering the scenario of a 200-km lunar orbit), the failure rate of a given storage device can be calculated using Eq. (5.12), as tabulated in TABLE 5-XVI. This table lists the spare capacity present in the device (commonly reported as the wear levelling percentage), and allows for 100,000 write cycles per cell. This then results in a total device failure (in which all cells have failed) after 117 years of continuous write operations, and a failure of all of the cells reserved for wear levelling after 10.6 years of continuous operation of the device. Note that this scenario assumes real-time RFI mitigation, as otherwise the raw write rate would amount to over 3 Gbit/s. Also, given the random nature of the data, it is not unlikely that the data will remain largely similar, which reduces the wear on the device.

TABLE 5-XVI
SSD LIFETIME PREDICTION BASED ON ALLOWED WRITE CYCLES

Device size	1024	GiByte
Wear levelling percentage	10	%
Write rate	0.263	Gbit/s
Number of write cycles allowed	100,000	-
Total device failure after	3.69E+09	seconds
	i.e. 117	years
10% failure after	3.35E+08	seconds
	i.e. 10.6	years

An OLFAR satellite can then be modelled according to the Markov²⁰ tree shown in Fig. 5.29, in a method identical to the one used for the models in sections 3.1.2, 3.1.3 and 3.1.4. In this case, no internal redundancy is assumed, and all soft-errors are assumed to be corrected for, as such an approach appears to be valid (see section 3.1.2).

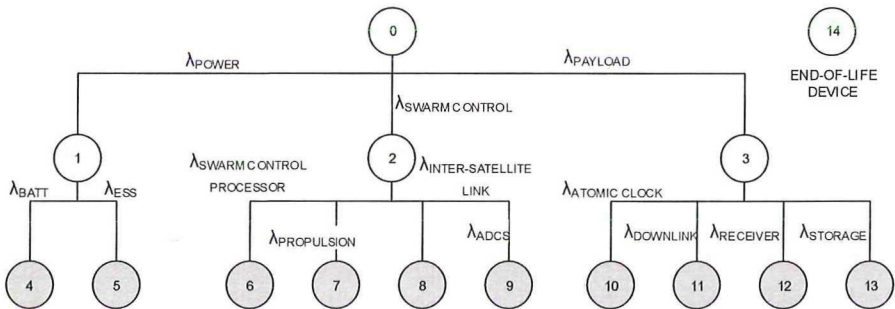


Fig. 5.29: Markov Tree for an OLFAR satellite

As can be seen in the figure, the end-of-life device is considered to be entirely independent, and therefore does not affect the reliability of the individual element. Each of the three primary functions groups their sub-functions, and no inter-relation between them is considered. This implies that, for example for the power system, when either the battery or the ESS fails, the power system is considered to fail permanently. This

²⁰ For a brief overview of the basic process of a Markov model, see Appdx A.

approach simplifies the analysis, yet can only be considered to be valid for an MTTFF analysis, in which case the first failure in any of the systems is dominant. This analysis is considered to be valid for an OLFAR element though, as without internal redundancy, the loss of any of the individual functionalities effectively renders the satellite useless as far as the swarm is considered (see also section 0), as taking over tasks of other defunct satellites actually decreases the reliability of the overall system.

The resulting MTTFF for an individual OLFAR element then equals 27.12 years, when using component lifetimes as in TABLE 5-XV, and a Markov model based on the overview in Fig. 5.29. An overview of the component lifetimes is shown in TABLE 5-XVII, together with the 90%-reliability point. Note these estimates assume an exponential failure distribution; which has been shown not to be realistic for complex systems. Weibull-distribution based Markov analysis is still a highly controversial area of research, although some progress has been made (Van Casteren, 2001). Since the reported shape-parameter for small satellites is 0.3134 (Monas, et al., 2012), infant mortality cases are the dominant failure mechanism. This implies that the estimates obtained for the OLFAR satellites are too optimistic for the initial part of the mission, yet on the conservative side once the period where all infant mortality cases occur has passed. This can, for the swarm, be solved through adding a number of extra elements; compensating for the infant mortality rate. This is a unique feature of a distributed system, as replacing faulty components in a monolithic satellite is troublesome at best.

TABLE 5-XVII
COMPUTED MTTFF OF AN OLFAR ELEMENT AND ITS SUB-FUNCTIONS

	MTTFF [year]	90% reliability based on MTTFF [year]
Energy supply	37.5	1.7
Swarm control	1668.6	76.4
Payload	104.1	4.8
Overall OLFAR element	27.1	2.8

5.4 PREDICTED PERFORMANCE

When using the OLFAR elements as designed in section 5.3 in a satellite swarm, the system lifetime can be computed, given a minimal number of 24 active elements, which was shown in section 5.2.1 to be able to achieve 49% of the required baseline coverage already in 100 days. The result of which is shown in Fig. 5.30. As the figure shows, the lifetime of the swarm can be extended up to over 43 years when expanding the swarm to 100 elements. When considering the 90% reliability-point, the lifetime of the swarm can be expanded from 2.8 years for the individual element to over four years.

The predicted element lifetime of 2.8 years, at which point the reliability of the element drops below 90%, is sufficient for the proposed 200-km lunar orbit scenarios, shown in TABLE 5-XII. A lifetime of 2.8 years is even sufficient for the 800-day 3000-km scenario, albeit barely. Expanding the number of elements in the swarm to 30 extends the useful lifetime to 3.3 years, allowing for more margin in meeting the 800-day requirement.

The system operational lifetime (i.e. availability) can also be simulated in a Monte Carlo analysis, focusing on the available observation time with a given system scenario. This allows predicting when the array would have to be replenished to allow uninterrupted

observations, or up to which point the array would be useful with a given initial set of satellites.

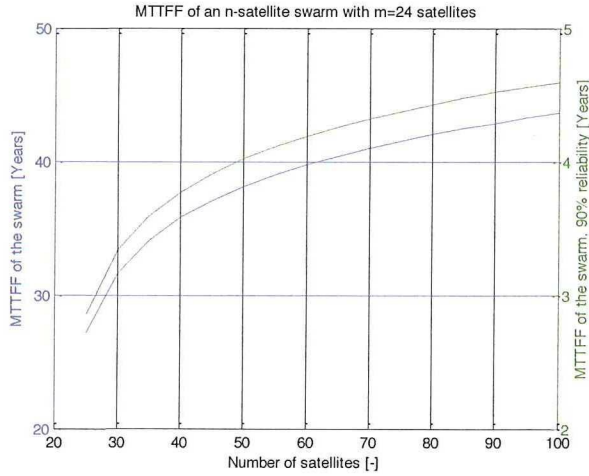


Fig. 5.30: MTTFF of an n-m swarm consisting of n OLFAR elements

This has been simulated for an initial array of 25 satellites, with input parameters as given in TABLE 5-XVIII, and the results are shown in Fig. 5.31 and Fig. 5.32, which display the available baselines over the mission duration and the accumulated observation time, respectively. The same can be repeated for 100 satellites, as shown in Fig. 5.31 and Fig. 5.33.

TABLE 5-XVIII
INPUT PARAMETERS TO THE MONTE CARLO ANALYSIS

Parameter	Value
Lifetime variation distribution	Weibull
Shape parameter	0.3134 ^(A)
Number of satellites in the swarm	25 or 100
Minimum required operational satellites	24
Single element nominal lifetime	2.8 years
Number of simulation runs	1,000,000

^A Data taken from (Monas, et al., 2012)

The simulations were run for 10^6 iterations, at which point the Weibull probability distribution becomes clearly visible in the graph showing the number of available baselines over time. For the 100-satellite swarm some ruggedness is still visible however, yet the available computational time was insufficient for running it for more iterations, which would smoothen the curve. This would also cause the 100-satellite swarm to achieve a similar final lifetime as a 25 satellite swarm, as this is to be expected due to the probability distribution used to spread the individual element lifetimes. The baselines can then be integrated to show the total accumulated observation time. As can be seen, the total accumulated observation time for the swarm at 200 km orbital altitude meets the 52,000 days requirement after approximately 0.68 years for the 25-satellite scenario, and after approximately 0.41 years for the 100-satellite scenario. The 3000 km orbit scenario suffers from the reduction in the available eclipse time, as shown clearly by the

accumulated observation time; yet benefits from an increase in available processing time of the data.




Fig. 5.31: MTTFF simulation of the number of active baselines over time for a swarm consisting of 25 satellites or 100 satellites.

Most importantly, the lifetimes of the satellites, in these cases using the 90% MTTFF reliability figure of 2.8 years per element can be seen to limit the useful lifetime of the swarm, as nearly all data is collected before most of the elements experience their first failure. In case this element design is used, and more observations are desired, it would be recommended to replenish the array after approximately 2.5 years; which would extend the mission lifetime by another 2.5 years.

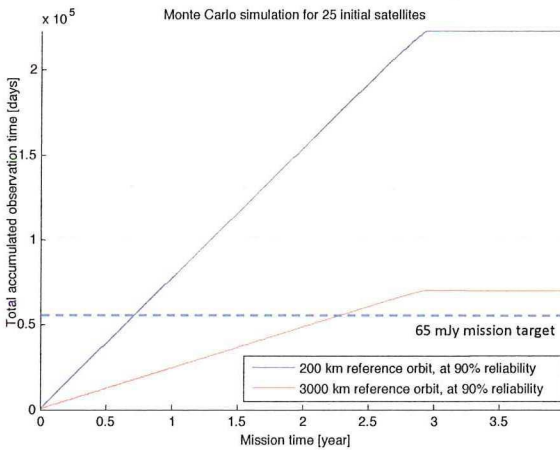


Fig. 5.32: MTTFF simulation of the total accumulated observation time (1) for a swarm with 25 satellites.




Fig. 5.33: MTTF simulation of the total accumulated observation time (2) for a swarm with 100 satellites.

Note that this simulation does not take single event upsets into account. These will result in a loss of baselines for observations, similar to losses due to RFI-events, as the elements experiencing upsets will have to recover from the event. The rate at which such events occur can only be predicted accurately once the payload processing chain is defined however; as it is highly dependent on the sensitivity of this chain to charged particles, as well as the target environment. Effectively however this will result in a loss of one or two elements participating in the observations. Since the number of baselines is defined by $N_{baselines} = N_{satellites}(N_{satellites} - 1)$, losing two elements will severely impact the number of active baselines. It is therefore recommendable to compensate for the SEU-rate through adding additional swarm elements to compensate for the average loss of satellites due to SEU upsets. In the example of two lost satellites per observation window, a 25-satellite swarm would then require 27 satellites in orbit. Reducing the sensitivity of the payload to single events is an alternative solution, yet it is quite likely the more expensive solution.

Near the end of the satellite lifetime (past the MTTFF), the satellite reliability lies below 73%, which could result in a significant reduction in useful observation periods and baselines. While it is quite likely that the swarm would still be operational at this point in time, it should not be designed for, as it will be difficult to guarantee useable output. In contrast to traditional satellites however, it would be acceptable for satellite swarm designs to assume MTTFF's below the 90% reliability point, as the redundancy offered by the swarm, and hence the graceful degradation, would offer the chance at performing observations with the remaining elements. The overall reliability of an OLFAR swarm, assuming an element lifetime of 2.8 years is shown in Fig. 5.34. It should be clear that the

100-satellite swarm, which benefits from a significant amount of redundancy, retains its initial reliability longer than the 25-satellite swarm. At about 3.1 years however, the Weibull failure distribution of the elements causes all elements to stop functioning, which causes the swarm reliability to vanish as well. Due to the significant redundancy however, this drop is much steeper.

Relying on graceful degradation beyond a given reliability figure should be discouraged. This reliability figure depends on the type of observations and the number of satellites present in the swarm. In case of OLFAR, the number can be read from Fig. 5.31, using the requirement of a minimum number of 24 operational satellites, which equates to 552 baselines. As can be seen in the figure, the 100-satellite swarm can guarantee more than 9312 baselines up to 2.86 years into the mission, whilst the 25-satellite swarm can only guarantee its original 552 baselines for up to 2.83 years, as it has less excess elements. The 100-satellite swarm at this point however, due to the steepness of the Weibull distribution, loses all of its elements almost simultaneously, resulting in a total loss of the system, yet it has gathered much more data up to this point, potentially making it more effective, provided all data can be transported and processed effectively. In reality however, failures are rarely this simultaneous, so it is likely that some outliers might still persist. Whether they can still provide useful data is another matter however.




Fig. 5.34: MTTF simulation of the reliability of an OLFAR swarm for 25 or 100 satellites.

6 CONCLUSIONS

Satellite swarms are an attractive topic for research, as the research field is both wide and deep. The wideness is due to the variety of aspects which satellite swarms and indeed many other robotic swarms have to take into account. The depth of the research field is due to the significant benefits various optimisation techniques can offer, as well as due to the complexity imposed by the interactions between the elements. In order to establish an unambiguous basis for the research, the author has developed a definition as to what a satellite swarm entails, which defines a swarm as “a space system consisting of many identical, egalitarian spacecraft, cooperating to achieve a common global goal”. A number of research questions were posed, yet they still merit a concise answer.

The work presented in this thesis focused on a select number of specific sub-topics, based on this definition. While this particular definition might not necessarily become universally accepted, it serves as a guideline allowing research on the topic of satellite swarms to focus to the specific type of distributed space architecture described by the definition. This definition is seen as the answer to the first research question of this thesis, which reads “*Which definition of a satellite swarm would be the best fit within the category of existing and planned distributed space architectures?*”

The answer to the second question, “*Which types of application areas would be best suited for satellite swarms?*” had to be broken down into one application area per type of satellite swarm, as we had found that there were sub-types identifiable within the concept of satellite swarms. Based on the above definition, three distinct types of satellite swarms were identified and a novel categorisation was introduced. The identified types of swarms were found to differ primarily in the applied control strategy, resulting in a different orbital distribution. The first type of satellite swarms identified are so-called “satellite clouds” which perform no orbital corrections at all, and are therefore freely drifting. Through their cooperation these could be very useful when performing in-situ sensing, or when doing serendipitous sampling (i.e. relying on the statistical chance that when having a sufficiently large number of sensors at arbitrary locations, at least one sensor will have picked up an event). An example would be the QB50-mission, or the PlanetLabs flock. The second type of swarms identified is referred to as foraging satellite swarms. This type of swarms seems to be the most applicable to planetary observation, as due to their limited control-rules, they will form a loosely defined network of satellites. An example of an existing system with many similarities to a foraging swarm would be the Iridium constellation. The Iridium constellation differs primarily in the fact that the location of each of the elements is pre-determined, and fixed, in order to allow routing of the data packets to the ground station. In case of a foraging swarm, each element would be aware of the shortest path towards the ground station, or at least be able to route the signal in a similar way, through for example a gossip-like protocol. The last type of satellite swarms identified are so-called “schooling swarms”, which resemble schools of fish, or flocks of birds. This type is most useful when applied in interferometry applications or similar applications where close, well-defined proximity is imperative. The main difference between a schooling swarm, such as the proposed OLFAR telescope, and a formation flight like Darwin or the Terrestrial Planet Finder (TPF) is that in case of a schooling swarm it does not matter which element is at which location, as long as each predefined location remains occupied throughout each observation period. This

implies interchangeability of the elements, which was identified as one of the most defining aspects of a satellite swarm, is also key to this type of architecture.

An answer to the first part of the third research question, “*How to design and optimise a satellite swarm such that it achieves a certain mission goal?*” can be given in that satellite swarms, like most systems, have to be designed with a specific mission goal in mind. The benefits offered by a satellite swarm to a select type of missions should be treated as a feature of a swarm, and it therefore has to be designed for, as much as any other subsystem of a satellite. This includes modelling the interactions of the elements, and the resulting effects on the mission as a whole, as well as the size of the swarm. The lack of flight data of satellite swarms, and indeed practical designs of satellite swarms, requires adequate modelling in order to allow sizing a given satellite swarm design. This can be done using a Markov chain, which can be used as a model of the individual satellite, and then later on modelling the satellite swarm as a k-out-of-m system. This, in turn, renders estimates of the lifetime and reliability of the individual elements, and the overall swarm. This data can then be used to establish how many elements are required in the swarm in order to achieve a given lifetime, or a given throughput of the swarm, as prediction of the number of active payloads or transceivers in the system is possible. The result of an analysis performed by the author on a generalised model of a satellite swarm is that, contrary to the much touted benefit of graceful degradation, which concerns the second part of the third research question (“*Which effects does including graceful degradation add?*”) in a swarm, does not imply that elements should take over functionalities which are defective in other satellites. This is also seen in natural swarms, which either eat or cast out defective members, as they appear concerned with the overall viability of the swarm. This, given that natural swarms have seen billions of design iterations, further reinforces the simulations which show that taking over functionalities in other satellites in fact reduces the reliability of the overall system. The concept of installing a device or function in a given satellite to separate defunct swarm satellites from the swarm would, in light of this discovery, seem logical, as it allows the satellite to leave the swarm gracefully.

6
Compared to traditional satellites, the average lifetimes of nano-satellites are indeed much lower, and the reported reliabilities of nano-satellites are to date still far below those of traditional micro- and larger satellites. Using nano-satellites in a satellite swarm would therefore not necessarily offer an immediate benefit, as the swarm should at the very least meet its mission success criteria, which, as is the case with nano-satellites, if the individual elements are not reliable enough, cannot be guaranteed. The low cost and launch masses of nano- and pico-satellites however allow for additional units, which in turn can increase the system reliability. This method is only valid up to the point where the cost of a larger, more traditionally constructed satellite equals the cost of the total sum of all envisaged nano-satellites. Adding additional units does offer a benefit to the swarm, irrespective of the type of elements used and depending on the reliability model used. This is easily proven for a Gaussian probability distribution model for the element reliability. Yet the simulations performed for this thesis show that it also holds for more realistic Weibull distributions.

An important finding is that for satellite swarms, redundancy can be shifted from internal redundancy (i.e. redundant internal components and systems) to the swarm itself, as each of the individual swarm spacecraft can be seen as a hot spare of any other satellite, provided that the satellites are entirely interchangeable.

This finding in turn has an effect on the design process of the individual swarm members, which concerns the fourth research question (“*How to design the swarm elements*”

which, when operated as a satellite swarm, ensure the resulting satellite swarm achieves a given mission goal?). The design process of a swarm satellite will differ from the design process of a singular monolithic satellite, as the swarm, as an entity, performs certain functions, which would normally be placed locally in the case of a monolithic satellite, on a global level. One of those is functions is redundancy, which aims at increasing the satellite reliability through adding multiple units of each of the critical components. Other potential candidates are energy storage and communication bandwidth provision: given a swarm with a large number of elements, the overall energy consumption can be distributed and optimised, allowing each of the satellites to contain smaller energy storage devices. Similar methods can be applied to payloads and communication devices through forming distributed interferometers or distributed phased arrays. This has a significant impact on the design procedure, and detailed system simulations will have to be performed in order to ensure the swarm's functionality meets the design specifications. System-wide swarm simulations therefore become an integral part of the element design process.

Increased autonomy for satellite swarms is widely regarded as beneficial. However, increasing the autonomy of a satellite swarm increases the complexity of the design procedure. Cooperation, especially for satellite swarms, is imperative, yet close cooperation also causes emergent behaviour, which can render predictions of the resulting system behaviour a complex task, which significantly increases the complexity of the design procedure. The extra time spent in increasing the level of autonomy however will pay back when operating the swarm, as local management is not troubled as much by communication delays and visibility-considerations. Increased autonomy is expected to increase stability and hence the system reliability when using local controllers.

Throughout this thesis, the OLFAR swarm was used as a reference case, and the fifth research question, "*Which element design would suit the OLFAR mission?*" attempted to address some of the biggest concerns regarding the viability of the OLFAR concept.

The OLFAR swarm aims at forming a low-frequency radio telescope, orbiting at a large distance from Earth to avoid the strong radio interference in its vicinity. The science case driving OLFAR would revolutionise our understanding of the early universe, yet is also extremely demanding, as the signal strengths involved are extremely weak, and long integration times are required. This in turn requires stable yet slowly changing orbits. The original candidate orbit, a low-lunar orbit, was shown to be changing too rapidly to allow for practical systems to achieve a sufficiently large inter-satellite communication bandwidth to compensate for the required reduction of the observation integration time. The OLFAR system benefits from a large number of nodes, and would therefore benefit significantly from using as cheap a platform as possible, hence the interest of the design team in using nano-satellites for the swarm nodes. The feasibility of using nano-satellites as a platform has not been disproven in this research, and storage and deployment of the scientific payload, as well as power generation have been prototyped and shown to be viable both in terms of storage density as well as performance. The inter-satellite links have also been proven to provide sufficient bandwidth (for modest relative speeds between the individual nodes), and relative position determination accuracy as well as timing accuracy are also deemed feasible. The reliability of the individual satellites, as well as the overall swarm is modelled based on experiences from current nano- and pico-satellites. It was shown that with the current reliability estimates, using the relatively low performance data of past and current nano- and pico-satellites, a three year swarm lifetime would be possible, which should suffice for OLFAR's primary purposes.

Key open items are the exact imaging and interferometry algorithms, as well as the expandability of the system, as for certain science cases, the number of satellites will have to increase in order to achieve the mission's demanding sensitivity requirements. Also the exact deployment location of the swarm is unknown, yet an Earth-Moon Lagrange orbit or a very High Earth Orbit appear to be viable candidates.

The sixth question, "*How to design the most basic swarm satellite?*" still lacks an answer, and will partly be treated and answered in section 6.1 which attempts to identify an archetype of the most basic swarm satellite, as the design process of a basic swarm satellite can only be identified once the most basic swarm satellite has been characterised. However, given the lack of satellite swarms which have already flown, the adequacy of this design process can only prove its merits when applied to many different swarm missions. Even when the definition of the most basic swarm satellite is known, no definitive answer will be readily available, as much of these design processes will undoubtedly involve, initially, at least some level of trial and error, and the subsequent build-up of experience from lessons learnt.

6.1 WHAT MAKES A SATELLITE A SWARM SATELLITE?

Following the definition of "A satellite swarm can be defined as a space system consisting of many identical, egalitarian spacecraft, cooperating to achieve a common global goal", swarm satellites should at least cooperate, in order to achieve their common goal.

Such a broad task description is also valid for other types of distributed systems. The fact they should be egalitarian stands out however, which implies that no diversification should be present. This does not imply that swarm satellites have to be exact copies of one another, but that they should at least functionally behave similarly, be able to perform the same tasks with equal adeptness, and also be able to perform part of the swarm coordination at any given point in time. This only holds while the satellite partakes in the swarm. It is by no means impossible to allow satellites to participate in a swarm only for a limited period of time, after which they are free to resume other activities.

The bare functional minimum a swarm satellite should contain is some means of inter-satellite communication to allow cooperation which would lead to the swarm achieving their common goal. This can, in principle, be routed through a ground station or other form of relaying system. In fact, the entire swarm control, as well as all intelligence of all of the individual satellites could theoretically take part in the ground segment, with the space segment acting as remote terminals, carrying sensors and actuators. The long and irregular communication delays will not easily favour this solution however.

Coordination of swarms can be done through stimergetic means, or using global incentives, mimicking natural swarms. Such methods benefit from having local swarm intelligence, which requires some form of local controlling agent inside each of the swarm's satellites.

For the three distinct types of satellite swarms identified in section 2.1, one can see that the satellite clouds are the simplest form, as they do not require orbit control. This implies that their (local) controllers can dispense with orbit maintenance and collision avoidance activities, which limits them to observation planning activities only. In case of QB50 for example, sampling is likely to be post-correlated, and each of the individual

satellites can remain in a permanent observation mode, also relinquishing this task from the swarm control component. This in turn results in a satellite which only contains its essential components, such as some form of energy supply, a payload and a communication device, which transmits the gathered data either to a (dedicated) ground station, or even permanently broadcasting it. Many early CubeSats, for example the Delfi-C3 satellite, would therefore be able to participate in a satellite cloud.

Ironically also passive satellites, such as NASA's ECHO 1 (McDougal, et al., 1972), which was a passive spherical reflector, would also fit this description, as their energy is supplied externally, and their mere presence acts as an instrument and communication device.

Satellites in a foraging swarm however require orbit maintenance. This requires some means of orbit control actuation, as well as some form of communication device which allows for maintaining a sufficient spreading between the elements of the swarm. Since they only perform relative manoeuvres in order to maximise inter-element spacing, this communication device could simply be a short-range beacon, which, when received by another satellite, is to be avoided or moved away from, or it could consist of position information transmitted through the ground-segment, which effectively creates a virtual component for the distance sensing device. Moreover, some constellation management, other than the relative spreading operation is to be present, consisting primarily of collision avoidance services. This could also be implemented in a virtual manner through the ground segment; yet all satellites still require active means of controlling their relative orbits. This sets them apart from the simplest of nano-satellites or even most CubeSats to date. Satellites which feature attitude control could use differential drag methods, allowing them to take part in a foraging swarm, as could of course satellites which feature tethers or other thrust-generating devices. Observation planning, while useful in avoiding gathering redundant data, is not strictly required. Satellites could be in a mode of permanent observation, using for example nadir-pointing (optical) instruments.

The most complex satellite swarms are the schooling swarms, as they perform (close) formation maintenance. They are the most likely to require inter-satellite links with moderate to high communication bandwidths, in order to control, determine and maintain their relative positions. They will, to that end, also require some means of relative position determination. A simplified collision avoidance short-range beacon, which would suffice for foraging swarms will not suffice for schooling swarms. Active propulsion is also a necessity. Observation planning is technically also not required for schooling swarms, though in almost all practical cases such a feature will be present.

The essential components for each type of satellite swarm have been collected in TABLE 6-I. Note that some of the components mentioned in TABLE 6-I could be virtualised components, indicated by a 'V'. Also components which are not strictly required are shown between brackets ("(...)").

TABLE 6-1
 ESSENTIAL COMPONENTS FOR SWARM SATELLITES,
 PARTICIPATING IN DIFFERENT TYPES OF SATELLITE SWARMS

	Satellite cloud	Foraging swarm	Schooling swarm
Elementary components (Energy supply, payload, communication device)	X	X	X
Inter-satellite Communication		V/(X)	X
Simple orbit maintenance (collision avoidance)		X	(X)
Precise orbit control			X
Observation planning			(V)/(X)
Swarm management controller		V/(X)	(V)/X
Relative position determination		(V)	X

As can be seen, there is not one particular distinguishing feature for a swarm satellite which would or could not be present in a traditional satellite. Like other satellites, and indeed other robotic or even natural swarms, it requires a means of communication and coordination and in the case of foraging and schooling swarms, some form of actuation, yet such features are not all that different from other satellites. Most of these means can be virtualised, i.e. executed in a digitalised environment executed, for example, in the ground segment. However, locomotion and communication devices (at least the physical layer) are devices which have to be present on the satellite platform. Relative position determination is only required for schooling swarms, which target well-defined relative positions between the elements. The lack of a uniquely distinguishing feature of satellite swarm satellites to classical satellites has an upside though, in that the experiences gained from designing, constructing and operating monolithic satellites will also largely be valid for satellite swarms. Only aspects related to coordination, such as predicting emerging behaviour are novelties which require further study. This also implies that, at this point in time, one would be able to construct and launch a swarm of satellites, provided that some measures are installed to handle faulty behaviour of individual elements, and that for this particular mission, the control rules are well-defined and well-understood.

6

The answer to sixth research question (“*How to design the most basic swarm satellite?*”) will therefore for a large part be identical to the question of “How to design the most basic classical satellite”, taking in mind however that for foraging or schooling swarms, orbit control is required, and that emergent behaviour for that particular swarm will have to be addressed. Note that not all emergent behaviour has a negative impact on the swarm, and in certain cases no unexpected behaviour ever emerges, yet it would be prudent to simulate and test the system prior to launch, ensuring it will not become an issue. In case the functionality of the swarm or the mission goals rely on (positive) emergent behaviour, these simulations will have been performed well ahead of finalising the spacecraft design, which in turn will also highlight potential causes for concern, if any.

6.2 OUTLOOK

This research forges into the large and novel research field of satellite swarms. The definition provided here can serve as a guideline for many research topics. Satellite swarms could prove to be very useful tools in e.g. wide-area, fast surveys, or in transient detecting systems. Selected interferometry missions, like OLFAR, will also benefit from a swarm architecture, as long as the elements are interchangeable.

While this research has shown that reliability analysis of satellite swarms using traditional methods is possible, more up-to-date and accurate input on actual component lifetimes is required in order to allow assessing the lifetimes of individual elements in a satellite swarm. The design method proposed in this research relies on this data to allow determining the number of satellites which should be present in a satellite swarm. The availability of up-to-date and relevant data therefore will allow for making a more accurate estimate, which in turn will result in a reduction in cost, since as with any design, less of a contingency margin will be required once more accurate data is available.

More than ever, for the design of satellite swarm systems detailed simulations, preferably using (representative) hardware in the loop will be required during the design stages. Standardisation of the individual swarm satellite elements and interfaces will undoubtedly prove beneficial, as such standardised platforms can significantly reduce the mission cost due to large volume production, as well as through providing statistically significant amounts of performance data for use in system simulations. Using standardised, generic components will impose an overhead, and quite likely result in less-than-optimal element designs. Yet the reduction in cost due to the increased production volumes, as well as the amount of (flight) data from similar missions could reduce the mission cost, the risks involved and probably also the time-to-flight. However, more detailed research into optimisation of the design method for satellite swarms, in particular the swarm elements, will have to be performed. In the end, only actual data obtained during missions can be used for a complete validation of the adequacy of the design approach used for that particular mission.

An OLFAR-like system, once operational, will provide invaluable insights into the early universe, as well as other possibly unknown phenomena occurring in the low frequency regime. Technologically speaking though, OLFAR still has a long road ahead. The primary candidate science orbit, which is a low lunar orbit, provides significant obstacles in terms of required inter-satellite bandwidth which has to be overcome before such an orbit would be viable. Lagrangian orbits could prove to be a more appropriate due to the reduced relative velocities, yet their remoteness is an issue. High Earth orbits however show exemplary performance in terms of relative velocities, and are much more accessible. Moreover, since they lack shielding by the moon, their data will contain much more Radio Frequency Interference. In turn however the satellites can continuously offload data to Earth, as well as receive clock corrections from an Earth-based atomic reference clock. Sending a precursor array to a high Earth orbit would then allow scientists to gather the first scientifically relevant data, whilst also allow mission designers and operators to gain experience in operating a satellite swarm in a relatively remote location. Should the data prove to be too noisy for the most sensitive science cases due to the Radio Frequency Interference (RFI) induced by Earth, the experience gained can then still be used to construct a new array in either an LL2 orbit, or even in an Earth-leading or trailing orbit around the Sun.

To date no technological barriers have been encountered which would disprove the viability of the OLFAR concept of using a swarm of nano-satellites in a remote orbit. The biggest technological hurdle on the way to realising OLFAR is to develop a distributable imaging algorithm. From a research perspective, except for the imaging algorithm, most obstacles appear to have been addressed, which could open the door towards its implementation.

REFERENCES

- Alexander, J. K., Kaiser, M. L., Novaco, J. C., Grena, F. R., & Weber, R. R. (1975). Scientific Instrumentation of the Radio-Astronomy-Explorer-2 Satellite. *Astronomy & Astrophysics* (40), 365-371.
- Apostolakis, G. E. (1976). The effect of a certain class of potential common mode failures on the reliability of redundant systems. *Nuclear Engineering and Design*, 36(1), 123-133.
- Armitage, S. E., Stras, L., Bonin, G. R., & Zee, R. E. (2013). The CanX-4/-5 Nanosatellite Mission And Technologies Enabling Formation Flight. *7th International Workshop on Satellite Constellations and Formation Flying*, Lisbon, Portugal.
- Asenek, V. A., Underwood, C. I., Oldfield, M. K., & Ward, J. W. (1997). Predicting the Influence of Software on the Reliability of Commercial off-the-shelf(COTS) Technology Microprocessors in a Space Radiation Environment. *11th ALAA/USU Conference on Small Satellites*, (pp. 1-8). Logan, Utah, US.
- Aso, S., Yokota, S., Hashimoto, H., Ohyama, Y., Sasaki, A., & Kobayashi, H. (2008). Control and Stability for Robotic Swarm based on Center of Gravity of Local Swarm. *IEEE International Symposium on Industrial Electronics*. Cambridge, United Kingdom.
- Autenrieth, A., & Kirstädter, A. (2000). Fault-tolerance and Resilience Issues in IP-based Networks. *Second International Workshop on the Design of Reliable Communication Networks (DRCN)*. München, Germany.
- Baan, W. A. (2012). *Space-based Ultra-long wavelength Radio Observatory – Low Cost Mission, ESA S-Class Mission Proposal*. ASTRON.
- Bandecchi, M., Melton, B., & Ongaro, F. (1999, September). Concurrent Engineering Applied to Space Mission Assessment and Design. *ESA Bulletin* (99).
- Barnhart, D. J. (2008). *Very Small Satellite Design for Space Sensor Networks*. Guilford, Surrey: University of Surrey.
- Bedingfield, K. L., Leach, R. D., & Alexander, M. B. (1996). *NASA Reference Publication 1390: Spacecraft System Failures and Anomalies Attributed to the Natural Space Environment*. National Aeronautics and Space Administration. Alabama, US: Marshall Space Flight Center - MSFC.
- Beekman, M., & Ratnieks, F. L. (2000). Long-range foraging by the honey-bee, *Apis mellifera* L. *Functional Ecology* (14), 490-496.
- Beichman, C. A., Woolf, N. J., & Lindensmith, C. A. (1999). *The Terrestrial Planet Finder (TPF) : a NASA Origins Program to search for habitable planets*. Washington, D.C., US: JPL.
- Belien, F., Engelen, S., Noroozi, A., & Verhoeven, C. (2011). Autonomous Navigation for Trans-lunar Nano-satellite Missions. *62nd International Astronautical Congress 2011*, (pp. 1-8). Cape Town, South Africa.

- Benjamin, J. L., & Paté-Cornell, M. E. (2004). Risk Chair for Concurrent Design Engineering: Satellite Swarm Illustration. *Journal of Spacecraft and Rockets*, 41(1), 51-59.
- Bennett, C. L., Larson, D., Weiland, J. L., Jarosik, N., Hinshaw, G., Odegard, N., Smith, K. M., Hill, R. S., Gold, B., Halpern, M., Komatsu, E., Nolta, M. R., Page, L., Spergel, D. N., Wollack, E., Dunkley, J., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., & Wright, E. L. (2013, September 20). Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *The Astrophysical Journal Supplement Series*, 208(2). doi:10.1088/0067-0049/208/2/20
- Bentum, M. J., & Boonstra, A.-J. (2009). OLFAR, Orbiting Low frequency Antennas for Radio Astronomy. *JENAM 2009, Enabling technologies for space-based astronomy and space*. London.
- Bentum, M. J., Verhoeven, C. J., Boonstra, A.-J., Van der Veen, A.-J., & Gill, E. K. (2009). A Novel Astronomical Application for Formation Flying Satellites. *60th International Astronautical Congress*. Daejeon, Republic of Korea.
- Bonnet, G., & Tessier, C. (2007). Collaboration among a satellite swarm. *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems (AAMAS '07)*. Honolulu, US.
- Boonstra, A.-J., Saks, N., Bentum, M. J., Falcke, H., Klein-Wolt, M., Rajan, R. T., Wijnholts, S. J., Arts, M., van 't Klooster, K., & Belien, F. (2010). DARIS - A Low-Frequency Distributed Aperture Array for Radio Astronomy in Space. *61th International Astronautical Congress*. Prague, Czech Republic.
- Boshuizen, C. R., Mason, J., Klupar, P., & Spanhake, S. (2014). Results from the Planet Labs Flock Constellation. *28th Annual ALAA/USU Conference on Small Satellites*, (pp. 1-8). Logan, Utah, USA.
- Brock, V. E., & Riffenburgh, R. H. (1960). Fish schooling: a possible factor in reducing predation. *Journal du Conseil*, 3(25), 307-317.
- Budianu, A., Meijerink, A., Bentum, M. J., Klein, J. M., & Engelen, S. (2014). Integrated Downlink Antennas in the Deployable Solar Panels of a Cubesat. *IEEE Aerospace Conference*, (pp. 1-11). Big Sky, Montana, US.
- Budianu, A., Rajan, R. T., Engelen, S., Meijerink, A., Verhoeven, C. J., & Bentum, M. J. (2011). OLFAR: Adaptive Topology for Satellite Swarms. *62nd International Astronautical Congress*. Cape Town, South Africa.
- Budianu, A., Willink Castro, T. J., Meijerink, A., & Bentum, M. J. (2013). Inter-satellite links for cubesats. *IEEE Aerospace Conference*, (pp. 1-10). Big Sky, Montana, US. doi:10.1109/AERO.2013.6496947
- Budianu, A., Willink-Castro, T. J., Meijerink, A., & Bentum, M. J. (2012). Communication Schemes for OLFAR's Inter-satellite Links. *63rd International Astronautical Congress*, (pp. 1-5). Naples, Italy.
- Buede, D. M. (2009). *The Engineering Design of Systems: Models and Methods*. Wiley.

- Buist, P. J., Noroozi, A., Sundaramoorthy, P., Verhagen, S., & Verhoeven, C. (2011, 6 1). Overview of Pulsar Navigation: Past, Present and Future Trends. *Navigation*, 58(2), 153-164.
- California Polytechnic State University. (2009). *Cubesat Design Specification, Rev. 12*. San Luis Obispo, California, US. Retrieved 2013, from http://www.cubesat.org/images/developers/cds_rev12.pdf
- California Polytechnic State University. (2013). *Cubesat Design Specification Rev. 13 (Provisional)*. San Luis Obispo, California, US. Retrieved 2013, from http://www.cubesat.org/images/developers/cds_rev13_draft_b.pdf
- Carmichael, C. H., & Brinkley, P. E. (2006, April 25). *Patent No. US 7036059B1*. United States of America.
- Carozzi, T. D., & Woan, G. (2009). A generalized measurement equation and van Cittert-Zernike theorem for wide-field radio astronomical interferometry. *Monthly Notices of the Royal Astronomical Society* (395), 1558-1568. doi:10.1111/j.1365-2966.2009.14642.x
- Carpenter, J. R., Folta, D. C., & Moreau, M. C. (2004). Libration point navigation concepts supporting the vision for space exploration. *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, (pp. 1-16). Providence, Rhode Island, US.
- Castet, J.-F., & Saleh, J. H. (2009, September-October). Satellite Reliability: Statistical Data Analysis and Modeling. *Journal of Spacecraft and Rockets*, 46(5), 1065-1076.
- Castet, J.-F., & Saleh, J. H. (2010). Beyond reliability, multi-state failure analysis of satellite subsystems: A statistical approach. *Reliability Engineering and System Safety* (95), 311-322.
- Castet, J.-F., & Saleh, J. H. (2012). On the concept of survivability, with application to spacecraft and space-based networks. *Reliability Engineering and System Safety* (99), 123-138.
- Chepuri, S. P., Rajan, R. T., Leus, G., & van der Veen, A.-J. (2013, January). Joint Clock Synchronization and Ranging: Asymmetrical Time-stamping and Passive Listening. *IEEE Signal Processing Letters*, 20(1), 51-54.
- Chin, A., Coelho, R., Brooks, L., Nugent, R., & Puig-Suari, J. (2008). Standardization Promotes Flexibility: A Review of CubeSats' Success. *6th Responsive Space Conference*, (pp. 1-9). Los Angeles, California, US.
- Chittka, L., & Niven, J. (2009, November 17). Are Bigger Brains Better? *Current Biology*, 19, R995-R1008. doi:10.1016/j.cub.2009.08.023
- Chou, A., Yang, J., Chelf, B., Hallem, S., & Engler, D. (2001). An empirical study of operating systems errors. *Proceedings of the eighteenth ACM symposium on Operating systems principles*, (pp. 73-88). New York, US. doi:10.1145/502034.502042
- Chu, J., Guo, J., & Gill, E. (2013). A survey of autonomous cooperation of modules' cluster operations for fractionated spacecraft. *International Journal of Space Science and Engineering*, 1(1), 3-19.

- Clement, B. J., & Barrett, A. C. (2002). Coordination Challenges for Autonomous Spacecraft. *AAMAS-02 workshop notes on Toward an Application Science: MAS Problem Spaces and Their Implications to Achieving Globally Coherent Behavior*.
- Colorado State University. (2013). Retrieved from Explorer-series reference images: http://rammb.cira.colostate.edu/dev/hillger/Explorer_reference.htm
- Curtis, S. A., Mica, J., Nuth, J., Marr, G., Rilee, M., & Bhat, M. (2000). ANTS (Autonomous Nano-Technology Swarm): An Artificial Intelligence Approach to Asteroid Belt Resource Exploration. *51st International Astronautical Federation Congress*. Rio de Janeiro, Brazil.
- Cuttler, J., Bennett, M., Klesh, A., Bahcivan, H., & Doe, R. (2010). The Radio Aurora Explorer – A Bistatic Radar Mission to Measure Space Weather Phenomenon. *The 24th Annual ALAA/USU Conference on Small Satellites*, (pp. 1-8). Logan, Utah, US.
- d'Amico, S., Ardaens, J.-S., De Florio, S., & Montenbruck, O. (2008). Autonomous Formation Flying - TanDEM-X, PRISMA and Beyond. *5th International Workshop on Satellite Constellations & Formation Flying*. Evpatoria, Crimea.
- Dekens, E. (2012). *Orbit Analysis of a Low Frequency Array for Radio Astronomy*. M.Sc. Thesis, Delft University of Technology, Space Engineering.
- Dekens, E. H., Engelen, S., & Noomen, R. (2014, September). A satellite swarm for radio astronomy. *Acta Astronautica, Vol. 102*, 321-331.
- Deng, J., & Shao, X.-M. (2006). Hydrodynamics in a diamond-shaped fish school. *Journal of Hydrodynamics, Ser. B, 18*(3, Supplement), 438-442.
- Dubos, G. F., Castet, J.-F., & Saleh, J. H. (2009). Statistical reliability analysis of satellites by mass category: does spacecraft size matter? *The 60th International Astronautical Congress*. Daejeon.
- Engelen, S., Gill, E., & Verhoeven, C. (2011). Systems Engineering Challenges for Satellite Swarms. *Aerospace Conference* (pp. 1-8). Big Sky, MT: IEEE. doi:10.1109/AERO.2011.5747259
- Engelen, S., Gill, E., & Verhoeven, C. (2012). On the reliability of spacecraft swarms. *Proceedings Of The 45 Symposium 2012*. Portoroz.
- Engelen, S., Gill, E., & Verhoeven, C. (2014, April). On the Reliability, Availability and Throughput of Satellite Swarms. *IEEE Transactions on Aerospace and Electronic Systems, Vol. 50*(2), 1027-1037.
- Engelen, S., Quillien, K. A., Verhoeven, C., Noroozi, A., Sundaramoorthy, P., van der Veen, A.-J., Rajan, R. T., Boonstra, A.-J., Bentum, M., Meijerink, A., & Budianu, A. (2013). The road to OLFAR-a roadmap to interferometric long-wavelength radio astronomy using miniaturized distributed space systems. *International Astronautical Congress* (pp. 1-7). Beijing, China: International Astronautical Federation.
- Engelen, S., van den Oever, M., Mahapatra, P., Sundaramoorthy, P. P., Gill, E. K., Meijer, R., & Verhoeven, C. J. (2012). NanoSAR – Case Study Of Synthetic Aperture Radar For Nano-Satellites. *63rd International Astronautical Congress*, (pp. 1-5). Naples, Italy.

- Engelen, S., Verhoeven, C., & Bentum, M. (2010). OLFAR, A Radio Telescope Based On Nano-Satellites In Moon Orbit. *24th Annual ALAA/USU Conference on Small Satellites*. Logan, UT.
- ESA. (2010, July 5). *The microwave sky as seen by Planck*. Retrieved from Space in Images: http://www.esa.int/spaceinimages/Images/2010/07/The_microwave_sky_as_seen_by_Planck
- Escorial, D., Tourne, I. F., Reina, F. J., Gonzalo, J., & Garrido, B. (2003). Fuego: A dedicated constellation of small satellites to detect and monitor forest fires. *Acta Astronautica*, 52(9-12), 765-776.
- Evison, S. E., Hart, A. G., & Jackson, D. E. (2008). Minor workers have a major role in the maintenance of leafcutter ant pheromone trails. *Animal Behaviour*(75), 963-969. doi:10.1016/j.anbehav.2007.07.013
- Ewen, H. I., & Purcell, E. M. (1951, September 1). Observation of a Line in the Galactic Radio Spectrum. *Nature*, 168, 356.
- Finn, A. M. (1989). System Effects Of Single Event Upsets. *Computers in Aerospace VII Conference* (pp. 994-1002). Monterey, CA, US: American Institute of Aeronautics and Astronautics, Inc.
- Fischer, G., & Rucker, H. O. (2006). Man-made Radio Emissions Recorded by Cassini/RPWS During Earth Flyby. (H. Rucker, W. Kurth, & G. Mann, Eds.) *Planetary Radio Emissions VI*, 299-306.
- Gill, E., d'Amico, S., & Montenbruck, O. (2007). Autonomous Formation Flying for the PRISMA Mission. *Journal of Spacecraft and Rockets*, 44(3), 671-681. doi:10.2514/1.23015
- Gill, E., Hamann, R., Monna, G. L., Scherpen, J. M., & Verhoeven, C. J. (2007). MISAT: Designing a Series of Powerful Small Satellites based upon Micro Systems Technology. *Proceedings of the 58th International Astronautical Congress (IAC)*. Hyderabad, India.
- Gill, E., Sundaramoorthy, P. P., Bouwmeester, J., Zandbergen, B., & Reinhard, R. (2013). Formation flying within a constellation of nano-satellites: The QB50 mission. *Acta Astronautica*, 82(1), 110-117.
- Ginosar, R. (2012). Survey of Processors For Space. *DASIA 2012*, (pp. 1-5). Dubrovnik.
- Gould, J. L. (1975). Honey Bee Recruitment: The Dance-Language Controversy. *Science*, 189, 685-693.
- Grabbe, C. L. (1981). Auroral Kilometric Radiation: A Theoretical Review. *Reviews of Geophysics and Space Physics*, 19(4), 627-633.
- Guo, J., Monas, L., & Gill, E. (2014). Statistical analysis and modelling of small satellite reliability. *Acta Astronautica*, 98, 97-110.
- Harrison, J. F., Kaiser, A., & VandenBrooks, J. M. (2010). Atmospheric oxygen level and the evolution of insect body size. *Proceedings of the Royal Society Biological Sciences*, 1-10. doi:10.1098/rspb.2010.0001
- Harriss, R. D., Holland, A. D., Barber, S. J., Karout, S., Burgon, R., Dryer, B. J., Murray, N. J., Hall, D. J., Smith, P. H., Grieg, T., Tutt, J. H., Endicott, J., Jerram, P.,

- Morris, D., Robbins, M., Prevost, V., & Holland, K. (2011). Compact CMOS Camera Demonstrator (C3D) for Ukube-1. *Proc. SPIE* 8146. doi:10.1117/12.895986
- Haslam, C. G., Salter, C. J., Stoffel, H., & Wilson, W. E. (1982). A 408 MHz All-Sky Continuum Survey. II. The Atlas of Contour Maps. *Astronomy & Astrophysics Supplement Series* (47), 1-143.
- Hegarty, C. J., & Chatre, E. (2008). Evolution of the Global Navigation Satellite System (GNSS). *Proceedings of the IEEE*, 96(12), 1902-1917.
- Heusdens, R., Engelen, S., Buist, P. J., Noroozi, A., & Sundaramoorthy, P. P. (2012). Match Filtering Approach for Signal Acquisition in Radio-Pulsar Navigation. *International Astronautical Conference* (pp. 1-5). Naples, Italy: International Astronautical Federation.
- Higashi, S., & Yamauchi, K. (1979). Influence of a Supercolonial Ant *Formica (Formica) Yessensis* Forel on the Distribution of Other Ants in Ishikari Coast. *Japanese Journal of Ecology*, 29, 257-264.
- Hinchey, M. G., Rash, J. L., Truszkowski, W. F., & Rouff, C. A. (2005). Autonomous and Autonomic Swarms. *International Conference on Software Engineering Research and Practice (SERP'05)*, (pp. 27-30). Las Vegas, Nevada, US.
- Hoevenaars, A. G., Engelen, S., & Bouwmeester, J. (2012). Model-Based Discrete PID Controller for CubeSat Reaction Wheels Based on COTS Brushless DC Motors. *Proceedings of the 1st LAA Conference on Dynamics and Control of Space Systems* (pp. 379-394). Porto, Portugal: Univelt, Inc.
- Hoff III, N. R., Sagoff, A., Wood, R. J., & Nagpal, R. (2010). Two Foraging Algorithms for Robot Swarms Using Only Local Communication. *IEEE International Conference on Robotics and Biomimetics*. Tianjin, China.
- Högbom, J. A. (1974). Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. *Astronomy and Astrophysics Supplement*, 15, 417-426.
- Hölldobler, B. (1976). The Behavioral Ecology of Mating in Harvester Ants (Hymenoptera: Formicidae: Pogonomyrmex). *Behavioral Ecology and Sociobiology* (1), 405-423.
- Huang, A., Chen, S.-S., Perng, H.-L., & Hsieh, M.-Y. (2009). Reliability allocation and prediction for developing small satellite. *7th LAA Symposium on Small Satellites for Earth Observation*. Berlin, Germany.
- International Space University. (2002). *Smaller Satellites: Bigger Business?* (M. Rycroft, & N. Crosby, Eds.) Springer Netherlands. doi:10.1007/978-94-017-3008-2
- Izzo, D., & Pettazzi, L. (2007). Autonomous and Distributed Motion Planning for Satellite Swarm. *Journal of Guidance, Control and Dynamics*, 30(2), 449-459.
- Jackson, B., & Epstein, K. (2000). A Reconfigurable Multifunctional Architecture Approach for Next-Generation Nanosatellite Design. *Proceedings of the IEEE Aerospace Conference*, 7, pp. 185-193. doi:10.1109/AERO.2000.879287
- Jafari, S., Ajorlou, A., Aghdam, A. G., & Tafazoli, S. (2010). Distributed Control of Formation Flying Spacecraft using Deterministic Communication Schedulers.

- 49th IEEE Conference on Decision and Control. Hilton Atlanta Hotel, Atlanta, GA, USA.
- Jähnichen, S., Briess, K., & Burmeister, R. (2008). Flying Sensors - Swarms in Space. In B. Mahr, & S. Huanye (Eds.), *Autonomous Systems – Self-Organization, Management, and Control* (pp. 71-77). Shanghai, China: Springer.
- Jansky, K. G. (1933). Electrical Disturbances Apparently of Extraterrestrial Origin. *Proceedings of the IRE*, 21(10), 1387-1398.
- Janson, S. W. (2011). 25 Years of Small Satellites. *25th Annual AIAA/USU Conference on Small Satellites*, (pp. 1-13). Logan, Utah, US.
- Jary, S. (2013, April 11). Least reliable smartphones revealed, Galaxy S3 most fault prone. *TechAdvisor*. Retrieved from <http://www.pcadvisor.co.uk/news/mobile-phone/3441592/least-reliable-smartphones-revealed/>
- Jester, S., & Falcke, H. (2009). Science with a lunar low-frequency array: From the dark ages of the Universe to nearby exoplanets. *New Astronomy Reviews* (53), 1-26.
- Johnson, R. A. (1991). Learning, Memory, and Foraging Efficiency in Two Species of Desert Seed-Harvester Ants. *Ecology*, 72(4), 1408-1419.
- Kahle, R., Schlepp, B., Aida, S., Kirschner, M., & Wermuth, M. (2012). Flight Dynamics Operations of the TanDEM-X Formation. *SpaceOps 2012*. Stockholm, Sweden.
- Kaiser, M. L., Desch, M. D., Bougeret, J.-L., Manning, R., & Meetre, C. A. (1996, May 15). Wind/WAVES observations of man-made radio transmissions. *Geophysical Research Letters*, 23(10), 1287-1290.
- Kardashev, N. S., Khartov, V. V., Abramov, V. V., Avdeev, V. Y., Alakoz, A. V., & al., e. (2013). "RadioAstron" - A Telescope with a Size of 300 000 km: Main Parameters and First Observational Results. *Astronomy Reports*, 57(3), 153-194.
- Kirkpatrick, D., Sellers, J. J., Thomas, D., & Verma, D. (2009). *Applied Space Systems Engineering*. (W. J. Larson, Ed.) McGraw-Hill Learning Solutions.
- Kirschner, M., Montenbruck, O., & Bettadpur, S. (2001). Flight Dynamics Aspects of the GRACE Formation Flying. *Proceedings of the 2nd International Workshop on Satellite Constellations and Formation Flying*. Haifa, Israel.
- Klein, J. M. (2014). *Design of OLFAR's energy supply*. M.Sc. Thesis, Delft University of Technology.
- Klein, J. M., Budianu, A., Bentum, M. J., Engelen, S., & Verhoeven, C. J. (2013). Design of an Electric Power System with Incorporation of a Phased Array Antenna for OLFAR. *The 64th International Astronautical Congress (IAC 2013)*, (pp. 1-6). Beijing, China.
- Klein-Wolt, M., Aminaei, A., Zarka, P., Schrader, J.-R., Boonstra, A.-J., & Falcke, H. (2012). Radio astronomy with the European Lunar Lander: Opening up the last unexplored frequency regime. *Planetary and Space Science* (74), 167-178.
- Klofas, B., & Leveque, K. (2013, April). *A Survey of CubeSat Communication Systems: 2009-2012*. Retrieved from http://www.klofas.com/papers/Klofas_Communications_Survey_2009-2012.pdf

- Kornienko, S., Kornienko, O., & Levi, P. (2004). Generation of desired emergent behavior in swarm of micro-robots. *ECAI*, 16, 239-243.
- Koulamas, C. (1992). Quality improvement through product redesign and the learning curve. *Omega*, 20(2), 161-168.
- Kramer, H. J., & Cracknell, A. P. (2008). An overview of small satellites in remote sensing. *International Journal of Remote Sensing*, 29(15), 4285-4337. doi:10.1080/01431160801914952
- Leitão, J., Pereira, J., & Rodrigues, L. (2007). HyParView: A Membership Protocol for Reliable Gossip-Based Broadcast. *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. Edinburgh, Scotland.
- Lejeune, S., & Warnant, R. (2008, April). A novel method for the quantitative assessment of the ionosphere effect on high accuracy GNSS applications, which require ambiguity resolution. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(6), 889-900.
- Levi, P., & Kernbach, S. (2010). Functional and Reliability Modelling of Swarm Robotic Systems. In *Symbiotic Multi-Robot Organisms* (pp. 5-77). Berlin Heidelberg: Springer-Verlag.
- Liang, H., Sun, Z., & Wang, J. (2013). Attitude coordination of satellite swarms with communication delays. *Aircraft Engineering and Aerospace Technology: An International Journal*, 85(3), 222-235. doi:10.1108/00022661311313713
- Liebe, H. J. (1983). Atmospheric EHF Window Transparencies near 35, 90, 140 and 220 GHz. *IEEE Transactions on Antennas and Propagation*, AP-31(1), 127-135.
- Liu, W., Winfield, A. F., Sa, J., Chen, J., & Dou, L. (2007). Towards Energy Optimization: Emergent Task Allocation in a Swarm of Foraging Robots. *Adaptive Behaviour*, 15(3), 289-305.
- Lovellette, M. N., Wood, K. S., Wood, D. L., Beall, J. H., Shirvani, P. P., Oh, N., & McCluskey, E. J. (2002). Strategies for Fault-Tolerant, Space-Based Computing: Lessons Learned from the ARGOS Testbed. *IEEE Aerospace Conference Proceedings*, 5, pp. 2109-2119. doi:10.1109/AERO.2002.1035377
- Lyke, J., Fronterhouse, D., Cannon, S., Lanza, D., & Byers, W. (2005). Space Plug-and-play Avionics. *ALAA 3rd Responsive Space Conference 2005*, (pp. 1-12). Los Angeles, California, US.
- Maccone, C. (2008). Protected antipode circle on the Farside of the Moon. *Acta Astronautica*, 63, 110-118.
- Maine, K., Devieux, C., & Swan, P. (1995). Overview of IRIDIUM satellite network. *WESCON/'95. Conference record. 'Microelectronics Communications Technology Producing Quality Products Mobile and Portable Power Emerging Technologies'*. San Francisco, US.
- Martin, Q., & George, A. D. (2012). Scrubbing optimization via availability prediction (SOAP) for reconfigurable space computing. *IEEE Conference on High Performance Extreme Computing (HPEC)*, 2012, (pp. 1-6). doi:10.1109/HPEC.2012.6408673

- Massioni, P., Keviczky, T., & Verhaegen, M. (2008). New Approaches to Distributed Control of Satellite Formation Flying. *The 3rd International Symposium on Formation Flying, Missions and Technologies*. ESA-ESTEC, Noordwijk, The Netherlands.
- McDougal, D. S., Lee III, R. B., & Romick, D. C. (1972). Measured Physical and Optical Properties of the Passive Geodetic Satellite (Pageos) and Echo 1. In S. W. Henriksen, *The Use of Artificial Satellites for Geodesy* (Vol. 15). American Geophysical Union.
- McEwen, J. D., & Wiaux, Y. (2011). Compressed sensing for wide-field radio interferometric imaging. *Monthly Notices of the Royal Astronomical Society* (413), 1318-1332. doi:10.1111/j.1365-2966.2011.18217.x
- Michener, C. D., & Grimaldi, D. A. (1988, September). The oldest fossil bee: Apoid history, evolutionary stasis, and antiquity of social behavior. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 85(17), 6424-6426.
- Miller, M. J. (2012, March 13). Mobile Core Wars: How the Chip Makers Stack Up. *PCMag*. Retrieved from <http://forwardthinking.pcmag.com/cell-phones/295297-mobile-core-wars-how-the-chip-makers-stack-up>
- Molette, P., Cougnet, C., Saint-Aubert, P., Young, R. W., & Helas, D. (1984). Technical and economical comparison between a modular geostationary space platform and a cluster of satellites. *Acta Astronautica*, 12(11), 771-784.
- Monas, L. N., Guo, J., & Gill, E. K. (2012). Small satellite reliability modeling: A statistical analysis. *The 4S Symposium 2012*. Portoroz.
- Montenbruck, O., Kahle, R., d'Amico, S., & Ardaens, J.-S. (2008). Navigation and control of the TanDEM-X formation. *The Journal of the Astronautical Sciences*, 56(3), 341-357.
- Moore, G. E. (1998, January). Cramming More Components onto Integrated Circuits. *Proceedings of the IEEE*, 86(1), 82-85.
- Morgan, D., Chung, S.-J., & Hadaegh, F. Y. (2013). Decentralized model predictive control of swarms of spacecraft using sequential convex programming. *23rd AAS/ALAA Space Flight Mechanics Meeting*. Kauai, Hawaii.
- Muller, C. A., & Oort, J. H. (1951, September 1). The Interstellar Hydrogen Line at 1,420 Mc./sec. and an Estimate of Galactic Rotation. *Nature*, 168(4270), 357-358.
- NASA / WMAP Science Team. (n.d.). *Nine Year Microwave Sky*. Retrieved from http://map.gsfc.nasa.gov/media/121238/ilc_9yr_moll4096.png
- Nicholson, W. L., Ricco, A. J., Agasid, E., Beasley, C., Diaz-Aguado, M., Ehrenfreund, P., Friedericks, C., Ghassemieh, S., Henschke, M., Hines, J. W., Kitts, C., Luzzi, E., Ly, D., Mai, N., Mancinelli, R., McIntyre, M., Minelli, G., Neumann, M., Parra, M., Piccini, M., Rasay, R. M., Ricks, R., Santos, O., Schooley, A., Squires, D., Timucin, L., Yost, B., & Young, A. (2011). The O/OREOS Mission: First Science Data from the Space Environment Survivability of Living Organisms (SESLO) Payload. *Astrobiology*, 11(10), 951-958. doi:10.1089/ast.2011.0714

- Nilsson, D. K., & Larson, U. E. (2008). Secure Firmware Updates over the Air in Intelligent Vehicles. *IEEE International Conference on Communications Workshops*, (pp. 380-384). Beijing. doi:10.1109/ICCW.2008.78
- Niven, J. E. (2012, January 6). How Honeybees Break a Decision-Making Deadlock. *Science*, 335, 43-44.
- Noca, M., Jordan, F., Steiner, N., Choueiri, T., George, F., Roethlisberger, G., Scheidegger, N., Peter-Contesse, H., Borgeaud, M., Krpoun, R., & Shea, H. (2009). Lessons Learned from the First Swiss Pico-Satellite: SwissCube. *23rd Annual ALAA/USU Conference on Small Satellites*, (pp. 1-20). Logan, Utah, US.
- Nolan, F. S., & Heap, H. F. (1978). *Reliability-Centered Maintenance*. United Airlines, San Francisco, US.
- Normann, R. D. (1998, May 19). *Patent No. 5,754,567*. United States of America.
- Novaco, J. C., & Brown, L. W. (1978). Nonthermal Galactic Emission Below 10 Megahertz. *The Astrophysical Journal* (221), 114-123.
- Nugent, R., Munakata, R., Chin, A., Coelho, R., & Puig-Suari, J. (2008). The CubeSat: The Picosatellite Standard for Research and Education. *ALAA SPACE 2008 Conference & Exposition*. San Diego, US.
- OHB-Sweden AB. (2012). *Innovative, Low Cost Satellites Mission Feasibility Report, Public Summary*. OHB-Sweden AB.
- Overlack, A. E., Kuiper, J. M., Peter-Contesse, H., & Noca, M. (2011). Analysis of the Attitude Control Instability of the SwissCube Nano Satellite. *Proceedings of the 1st IAA Conference on University Satellites Missions*. Roma, Italy.
- Parrish, J. K., Viscido, S. V., & Grünbaum, D. (2002). Self-Organized Fish Schools: An Examination of Emergent Properties. *the Biological Bulletin*, 202(3), 296-305.
- Partridge, B. L., Johansson, J., & Kalish, J. (1983). The structure of schools of giant bluefin tuna in Cape Cod Bay. *Environmental Biology of Fishes*, 9(3/4), 253-262.
- Paruchuri, V. K., Durresti, A., Durga, D. S., & Raj, J. (2003). Optimal Flooding Protocol For Routing In Ad-hoc Networks. *IEEE Wireless Communications and Networking Conference*. New Orleans, US.
- Paté-Cornell, M. E., Dillon, R. L., & Guikema, S. D. (2004). On the Limitations of Redundancies in the Improvement of System Reliability. *Risk Analysis*, 24(6), 1423-1436.
- PC/104 Consortium. (2013, September 13). *SPECIFICATIONS — PC/104*. Retrieved from http://www.pc104.org/pc104_specs.php
- Pincirolì, C., Birattari, M., Dorigo, M., del Rey Zapatero, M., Vinko, T., & Izzo, D. (2008). Lattice Formation in Space for a Swarm of Pico Satellites. In *Ant Colony Optimization and Swarm Intelligence* (pp. 347-354). Springer Berlin Heidelberg. doi:10.1007/978-3-540-87527-7_36
- Pincirolì, C., Birattari, M., Tuci, E., Dorigo, M., del Rey Zapatero, M., Vinko, T., & Izzo, D. (2008). *Evolving a Collective Consciousness for a Swarm of Pico-Satellites*. ESA ACT.

- Platt, T., Fuentes-Yaco, C., & Frank, K. T. (2003). Marine ecology: Spring algal bloom and larval fish survival. *Nature*, 423, 398-399.
- Prior, M. (2006). *IMAGE Failure Review Board Final Report*.
- Pukite, J., & Pukite, P. (1998). *Modeling for Reliability Analysis: Markov Modeling for Reliability, Maintainability, Safety, and Supportability Analyses of Complex Systems, 1st Edition*. Wiley-IEEE Press.
- Purcell, J. E., Brown, E. D., Stokesbury, K. D., Haldorson, L. H., & Shirley, T. C. (2000). Aggregations of the jellyfish *Aurelia labiata*: abundance, distribution, association with age-0 walleye pollock, and behaviors promoting aggregation in Prince William Sound, Alaska, USA. *Marine Ecology Progress Series*, 195, 145-158.
- Qin, F., Lu, S., & Zhou, Y. (2005). SafeMem: Exploiting ECC-Memory for Detecting Memory Leaks and Memory Corruption During Production Runs. *Proceedings of the 11th Int'l Symposium on High-Performance Computer Architecture (HPCA-11 2005)* (pp. 1-12). IEEE.
- Qiu, D., & Srikant, R. (2004). Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks. *SIGCOMM'04* (pp. 367-377). Portland, Oregon, US: ACM.
- Quillien, K. A. (2013). *An Astronomical Antenna for OLFAR*. M.Sc. Thesis, Delft University of Technology.
- Quillien, K. A., Engelen, S., Gill, E. K., Smith, D. M., Arts, M. J., & Boonstra, A.-J. (2013). Astronomical Antenna for a Space Based Low Frequency Radio Telescope. *27th Annual ALAA/USU Conference on Small Satellites*, (pp. 1-7). Logan, Utah, US.
- Rabbia, Y. (2004). Nulling Interferometry: from the Bracewell scheme to the ESA-Darwin mission. *EAS Publications Series*, 12, 215-234. doi:10.1051/eas:2004035
- Rajan, R. T., & van der Veen, A.-J. (2013). Joint non-linear ranging and affine clock synchronization basis for a network of mobile nodes. *EURASIP EUSIPCO*. Marrakech, Morocco.
- Rajan, R. T., Bentum, M. J., & Boonstra, A.-J. (2013). Synchronization for space based ultra low frequency interferometry. *IEEE Aerospace Conference*. Big Sky, Montana, US.
- Rajan, R. T., Bentum, M. J., Gunst, A., & Boonstra, A.-J. (2013). Distributed correlators for Interferometry in space. *IEEE Aerospace Conference*. Big Sky, Montana, US.
- Rajan, R. T., Engelen, S., Bentum, M., & Verhoeven, C. (2011). Orbiting Low Frequency Array for Radio Astronomy. *IEEE Aerospace Conference*. Big Sky, MT, US. doi:10.1109/AERO.2011.5747222
- Rajan, R. T., Leus, G., & van der Veen, A.-J. (2013). Relative Velocity Estimation Using Multidimensional Scaling. *IEEE CAMSAP*. St. Martin.
- Rinaldi, S. M., Peerenboom, J. P., & Kelly, T. K. (2001, December). Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies. *IEEE Control Systems Magazine*, pp. 11-25.

- Rouff, C., Vanderbilt, A., Hinchey, M., Truszkowski, W., & Rash, J. (2004). Verification of Emergent Behaviors in Swarm-based Systems. *11th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems*. doi:10.1109/ECBS.2004.1316730
- Rougoor, G. W., & Oort, H. J. (1960, January 15). Distribution and Motion of Interstellar Hydrogen in the Galactic System with Particular Reference to the Region within 3 Kiloparsecs of the Center. *Proceedings of the National Academy of Sciences*, 46(1), 1-13.
- Rubenstein, M., Ahler, C., & Nagpal, R. (2012). Kilobot: A Low Cost Scalable Robot System for Collective Behaviors. *IEEE International Conference on Robotics and Automation (ICRA)*, (pp. 1-6). St. Paul, Minneapolis, US.
- Rutishauser, S., Correll, N., & Martinoli, A. (2009). Collaborative coverage using a swarm of networked miniature robots. *Robotics and Autonomous Systems*, 57(5), 517-525. doi:10.1016/j.robot.2008.10.023
- Sabatini, M., & Palmerini, G. B. (2009). Collective control of spacecraft swarms for space exploration. *Celest. Mech. Dyn. Astro.*(105), 229-244. doi:10.1007/s10569-009-9183-8
- Sabol, C., Burns, R., & McLaughlin, C. A. (2001). Satellite Formation Flying Design and Evolution. *Journal of Spacecraft and Rockets*, 38(2), 270-278.
- Şahín, E. (2005). Swarm Robotics: From Sources of Inspiration to Domains of Application. In *Lecture Notes in Computer Science* (Vol. 3342, pp. 10-20). doi:10.1007/978-3-540-30552-1_2
- Saks, N., Boonstra, A.-J., Rajan, R. T., Bentum, M. J., Beliën, F., & van 't Klooster, K. (2010). DARIS, a fleet of passive formation flying small satellites for low frequency radio astronomy. *The 4S Symposium*. Madeira, Portugal.
- Saleh, J. H., Hastings, D. E., & Newman, D. J. (2004). Weaving time into system architecture: satellite cost per operational day and optimal design lifetime. *Acta Astronautica* (54), 413-431.
- Sánchez-Maestro, R., Agenjo Díaz, A., Escorial, D., Tarabini Castellani, L., & Cropp, A. (2013). PROBA-3 Control System for Fine Formation Flying. *7th International Workshop on Satellite Constellations and Formation Flying*. Lisbon, Portugal.
- Scherzer, O. (Ed.). (2011). *Handbook of Mathematical Methods in Imaging*. Springer New York. doi:10.1007/978-0-387-92920-0
- Schoeberl, M. R. (2002). The afternoon constellation: a formation of Earth observing systems for the atmosphere and hydrosphere. *IEEE International Geoscience and Remote Sensing Symposium (IGARSS '02)*.
- Seeley, T. D. (1983). Division of Labor Between Scouts and Recruits in Honeybee Foraging. *Behavioral Ecology and Sociobiology* (12), 253-259.
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marschall, J. A. (2012, January 6). Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms. *Science*, 335, 108-111.

- Shiroma, W. A., Martin, L. K., Akagi, J. M., Akagi, J. T., Wolfe, B. L., Fewell, B. A., & Ohta, A. T. (2011). CubeSats: A Bright Future for Nanosatellites. *Central European Journal of Engineering*, 1(1), 9-15. doi:10.2478/s13531-011-0007-8
- Spoelstra, T. A. (1997). The ionosphere and radio interferometry. *Annali di Geofisica*, XL(4), 865-885.
- Sunamura, E., Espadaler, X., Sakamoto, H., Suzuki, S., Terayama, M., & Tatsuki, S. (2009). Intercontinental union of Argentine ants: behavioral relationships among introduced populations in Europe, North America, and Asia. *Insectes Sociaux*, 56, 143-147. doi:10.1007/s00040-009-0001-9
- Sverdrup, H. U. (1953). On Conditions for the Vernal Blooming of Phytoplankton. *J. Cons. int. Explor. Mer*, 18(3), 287-295.
- Swartwout, M. (2012). A statistical survey of rideshares (and attack of the CubeSats, part deux). *IEEE Aerospace Conference* (pp. 1-7). Big Sky, MT, US: IEEE. doi:10.1109/AERO.2012.6187008
- Swartwout, M. A. (2011). A brief history of rideshares (and attack of the CubeSats). *IEEE Aerospace Conference* (pp. 1-15). Big Sky, MT, US: IEEE. doi:10.1109/AERO.2011.5747233
- Sweeting, M. N. (1992). UoSAT microsatellite missions. *Electronics & Communication Engineering Journal*, 141-150.
- Takahashi, Y. D. (2003). *New astronomy from the moon: a lunar based very low frequency radio array*. Master's thesis, University of Glasgow. Retrieved from <http://www.astro.gla.ac.uk/users/yuki/thesis/YukiThesis.pdf>
- Taraba, M., Rayburn, C., Tsuda, A., & MacGillivray, C. S. (2009). Boeing's Cubesat Testbed 1 Attitude Determination Design and On-Orbit Experience. *23rd Annual AIAA/USU Conference on Small Satellites*, (pp. 1-9). Logan, Utah, US.
- Tasse, C., van Diepen, G., van der Tol, S., van Weeren, R. J., van Zwieten, J. E., Batejat, F., Bhatnagar, S., van Bommel, I., Birzan, L., Bonafede, A., Conway, J., Ferrari, C., De Gasperin, F., Golap, K., Heald, G., Jackson, N., Macario, G., McKean, J., Mohan, N., Orrù, E., Pizzo, R., Rafferty, D., Rau, U., Röttgering, H., & Shulevski, A. (2012). LOFAR calibration and wide-field imaging. *Comptes Rendus Physique*, 13(1), 28-32.
- Thatcher, J., Coughlin, T., Handy, J., & Ekker, N. (2009). *NAND Flash Solid State Storage for the Enterprise, An In-depth Look at Reliability*. White paper, SNIA: Solid State Storage Initiative.
- Thompson, A. R., Moran, J. M., & Swenson Jr., G. W. (2001). *Interferometry and Synthesis in Radio Astronomy*, 2nd edition. Wiley.
- Thorne, B. L., Grimaldi, D. A., & Krishna, K. (2000). Early Fossil History of the Termites. In T. Abe, D. E. Bignell, & M. Higashi, *Termites: Evolution, Sociality, Symbioses, Ecology, Part 2* (pp. 77-93). Kluwer Academic Publishers.
- Tripp, H., & Palmer, P. (2010). Stigmergy based behavioural coordination for satellite clusters. *Acta Astronautica*, 66, 1052-1071.

- Truszkowski, W., Hinchey, M., Rash, J., & Rouff, C. (2004). NASA's swarm missions: The challenge of building autonomous software. *IT professional*, 6(5), 47-52.
- Tsinas, A., & Welch, C. (2001). Efficient Satellite Structural Design Optimised for Volume Production. *Proceedings of the 15th Annual USU Conference on Small Satellites*, (pp. 1-20). Logan, Utah, US.
- Ubbels, W. J., Bonnema, A. R., van Breukelen, E. D., Doorn, J. H., van den Eikhoff, R., Van der Linden, E., Aalbers, G. T., Rotteveel, J., Hamann, R. J., & Verhoeven, C. J. (2005). Delfi-C3: a Student Nanosatellite as a Test-bed for Thin Film Solar Cells and Wireless Onboard Communication. *Proceedings of the 2nd International Conference on Recent Advances in Space Tech*, (pp. 167-172). Istanbul, Turkey. doi:10.1109/RAST.2005.1512556
- Van Casteren, J. (2001). *Power System Reliability Assessment using the Weibull-Markov Model*. PhD. Thesis, Chalmers University of Technology, Department of Electric Power Engineering, Göteborg, Sweden.
- van Cittert, P. H. (1934). Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene. *Physica*, 1(1-6), 201-210.
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., Heald, G., McKean, J. P., & al., e. (2013). LOFAR: The LOw-Frequency ARray. *Astronomy & Astrophysics*, 556(A2), 1-53. doi:10.1051/0004-6361/201220873
- Verberk, W. C., & Bilton, D. T. (2011). Can Oxygen Set Thermal Limits in an Insect and Drive Gigantism? *PLoS ONE*, 6(7), 1-6. doi: doi:10.1371/journal.pone.0022610
- Verhoeven, C. J., & Jongkind, W. (2003). MISAT: A Satellite Colony. *8th International Conference on the Commercialization of Micro and Nano Systems*. Amsterdam, The Netherlands.
- Verhoeven, C., Bentum, M., Monna, B., Rotteveel, J., & Guo, J. (2011). On the origin of satellite swarms. *Acta Astronautica*, 68(7-8), 1392-1395.
- Verhoeven, C., Engelen, S., Noroozi, A., Sundaramoorthy, P., Bentum, M., & Meijer, R. (2011). Nanosatellite swarm support for larger satellites. *International Astronautical Congress*, (pp. 1-5). Cape Town, South Africa.
- Verhoeven, C., Monna, B., Engelen, S., Noroozi, A., Sundaramoorthy, P., Bentum, M., & Meijer, R. (2011). Emerging Eco-system: Nano-satellite Swarms and Large Satellites. *IAC 2011, IAC-11, D1,1,6,x11045*. Cape Town, South Africa.
- Vos, E., Shaik, F., Scherpen, J. M., & van der Schaft, A. J. (2013). Control For Formation Flying Of Satellites: Port-Hamiltonian Approach. *7th International Workshop on Satellite Constellations and Formation Flying*. Lisbon, Portugal.
- Weber, R. R., Alexander, J. K., & Stone, R. G. (1971, December). The Radio Astronomy Explorer satellite, a low-frequency observatory. *Radio Science*, 6(12), 1085-1097.
- Wenger, S., & Magnor, M. (2010). *SparseRI: A Compressed Sensing Framework for Aperture Synthesis Imaging in Radio Astronomy*. Technical Report, TU Braunschweig, Computer Graphics Lab. Retrieved from <http://www.digibib.tu-bs.de/?docid=00032047>

- Wertz, J. R., Everett, D. F., & Puschell, J. J. (2011). *Space Mission Engineering: The New SMAD*. Microcosm Press.
- Wilson, E. O., & Hölldobler, B. (2005, March 18). The rise of the ants: A phylogenetic and ecological explanation. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 102(21), 7411-7414. doi:10.1073/pnas.0502264102
- Wilson, E. O., Carpenter, F. M., & Brown Jr., W. L. (1967, September 1). The First Mesozoic Ants. *Science*, 157(3792), 1038-1040. doi:10.1126/science.157.3792.1038
- Wilson, E. O., Carpenter, F. M., & Brown, Jr., W. L. (1967, March). The First Mesozoic Ants, With the Description of a New Subfamily. *Psyche*, 74(1), 1-19.
- Wilson, T. L., Rohlf, K., & Hüttemeister, S. (2009). *Tools of Radio Astronomy, Fifth Edition*. Springer-Verlag Berlin Heidelberg.
- Winfield, A., Harper, C., & Nembrini, J. (2005). Towards dependable swarms and a new discipline of swarm engineering. In E. Şahin, & W. Spears (Ed.), *Swarm Robotics Workshop: State-of-the-art Survey* (pp. 126-142). Springer-Verlag.
- Winfield, A., Sa, J., Fernandez-Gago, M., Dixon, C., & Fisher, M. (2005). On formal specification of emergent behaviours in swarm robotic systems. *International Journal of Advanced Robotic Systems*, 2(4), 363-370.
- Winner, R. I., Pennel, J. P., Bertrand, H. E., & Slusarezuk, M. M. (1988). *The role of concurrent engineering in weapons system acquisition*. Alexandria: Institute for Defence.
- Winternitz, L. M., Bamford, W. A., & Heckler, G. W. (2009, August). A GPS Receiver for High-Altitude Satellite Navigation. *IEEE Journal of Selected Topics in Signal Processing*, 3(4), 541-556.
- Xue, S., & Zeng, J. (2008). Sense Limitedly, Interact Locally: the Control Strategy for Swarm Robots Search. *IEEE International Conference on Networking, Sensing and Control*. Hainan, China.
- Yashar, M., & Kembal, A. (2009). *Computational costs of radio imaging algorithms dealing with the non-coplanar baselines effect : I*. TDP Calibration & Processing Group (CPG) Memo.
- Zarka, P., Bougeret, J.-L., Briand, C., Cecconi, B., Falcke, H., Girard, J., Grießmeier, J.-M., Hess, S., Klein-Wolt, M., Konovalenko, A., Lamy, L., Mimoun, D., & Aminaei, A. (2012). Planetary and exoplanetary low frequency radio observations from the Moon. *Planetary and Space Science* (74), 156-166.
- Zavodnik, D. (1987). Spatial aggregations of the swarming jellyfish *Pelagia noctiluca* (Scyphozoa). *Marine Biology*, 94, 265-269.
- Zernike, F. (1938). The concept of degree of coherence and its application to optical problems. *Physica*, 5(8), 785-795.
- Zhang, D., Xie, G., Yu, J., & Wang, L. (2007). Adaptive task assignment for multiple mobile robots via swarm intelligence approach. *Robotics and Autonomous Systems*, 55, 572-588. doi:10.1016/j.robot.2007.01.008

APPENDICES

APPX. A: MEAN TIME TO FAILURE ANALYSIS USING MARKOV MODELLING

In order to analyse the availability and the reliability of satellites, a mean-time-to-failure analysis can be performed. This analysis renders the operational lifetime of a satellite, which can be used to assess the lifetime of a system of satellites in an analogous manner.

Determining the Mean Time To Failure (MTTF) is traditionally performed using a Markov Chain analysis (Pukite & Pukite, 1998). 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 a 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 1, 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 state 1 represents the state where a single subsystem has failed. States two and three then represent a state where another subsystem (subsystems two or three) has failed, after the first subsystem in this chain had already failed. Furthermore, λ_1 , λ_2 and λ_3 represent the failure rates between the different states, in units of failures per unit of (operating) time.




Figure 1: An elementary branch of a Markov Chain

Then this branch can be represented by the set of partial differential equations:

$$\frac{\partial P_0}{\partial t} = -\lambda_1 P_1 \quad (\text{A.1})$$

$$\frac{\partial P_1}{\partial t} = \lambda_1 P_0 - (\lambda_2 + \lambda_3) P_1 \quad (\text{A.2})$$

$$\frac{\partial P_2}{\partial t} = \lambda_2 P_1, \quad \frac{\partial P_3}{\partial t} = \lambda_3 P_1 \quad (\text{A.3})$$

in which P_i represents the propability of being in state 1.

Then, taking $\lim_{t \rightarrow \infty} P_1$ will result in the time taken for the system to arrive at P_1 . This modelling 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.




Figure 2: An elementary branch of a Markov Chain, including repairs between state 1 and 0

Repairs of a (sub-)system can be represented by an upward branch, with a repair-rate μ_n for the state 'n', and they can be subtracted from the differential equations according to:

$$\frac{\partial P_0}{\partial t} = -\lambda_1 P_1 + \mu_1 P_0 \quad (\text{A.4})$$

$$\frac{\partial P_1}{\partial t} = \lambda_1 P_0 - (\lambda_2 + \lambda_3) P_1 - \mu_1 P_0 \quad (\text{A.5})$$

$$\frac{\partial P_2}{\partial t} = \lambda_2 P_1, \quad \frac{\partial P_3}{\partial t} = \lambda_3 P_1 \quad (\text{A.6})$$

In this example, repairs between state 1 and 2 are allowed, whilst repairing the second and third subsystem are assumed to not be possible, and are hence not shown.

APPX. B: LIST OF PUBLICATIONS BY THE AUTHOR

A. AS PRIMARY AUTHOR

Engelen, S., Gill, E., & Verhoeven, C. (2011). Systems Engineering Challenges for Satellite Swarms. *Aerospace Conference* (pp. 1-8). Big Sky, MT: IEEE. doi:10.1109/AERO.2011.5747259

Engelen, S., Gill, E., & Verhoeven, C. (2012). On the reliability of spacecraft swarms. *Proceedings Of The 4S Symposium 2012*. Portoroz.

Engelen, S., Gill, E., & Verhoeven, C. (2014, April). On the Reliability, Availability and Throughput of Satellite Swarms. *IEEE Transactions on Aerospace and Electronic Systems, Vol. 50(2)*, 1027-1037.

Engelen, S., Verhoeven, C., & Bentum, M. (2010). OLFAR, A Radio Telescope Based On Nano-Satellites In Moon Orbit. *24th Annual AIAA/USU Conference on Small Satellites*. Logan, UT.

Engelen, S., Quillien, K. A., Verhoeven, C., Noroozi, A., Sundaramoorthy, P., van der Veen, A.-J., Rajan, R. T., Boonstra, A.-J., Bentum, M., Meijerink, A., & Budianu, A. (2013). The road to OLFAR-a roadmap to interferometric long-wavelength radio astronomy using miniaturized distributed space systems. *International Astronautical Congress* (pp. 1-7). Beijing, China: International Astronautical Federation.

Engelen, S., van den Oever, M., Mahapatra, P., Sundaramoorthy, P. P., Gill, E. K., Meijer, R., & Verhoeven, C. J. (2012). NanoSAR – Case Study Of Synthetic Aperture Radar For Nano-Satellites. *63rd International Astronautical Congress*, (pp. 1-5). Naples, Italy.

B. AS CO-AUTHOR

Beliën, F., Engelen, S., Noroozi, A., & Verhoeven, C. (2011). Autonomous Navigation for Trans-lunar Nano-satellite Missions. *62nd International Astronautical Congress 2011*, (pp. 1-8). Cape Town, South Africa.

Budianu, A., Meijerink, A., Bentum, M. J., Klein, J. M., & Engelen, S. (2014). Integrated Downlink Antennas in the Deployable Solar Panels of a Cubesat. *IEEE Aerospace Conference*, (pp. 1-11). Big Sky, Montana, US.

Budianu, A., Rajan, R. T., Engelen, S., Meijerink, A., Verhoeven, C. J., & Bentum, M. J. (2011). OLFAR: Adaptive Topology for Satellite Swarms. *62nd International Astronautical Congress*. Cape Town, South Africa.

Buist, P. J., Noroozi, A., Sundaramoorthy, P., Verhagen, S., & Verhoeven, C. (2011, 6 1). Overview of Pulsar Navigation: Past, Present and Future Trends. *Navigation*, 58(2), 153-164.

Dekens, E. H., Engelen, S., & Noomen, R. (2014, September). A satellite swarm for radio astronomy. *Acta Astronautica, Vol. 102*, 321-331.

Heusdens, R., Engelen, S., Buist, P. J., Noroozi, A., & Sundaramoorthy, P. P. (2012). Match Filtering Approach for Signal Acquisition in Radio-Pulsar

Navigation. *International Astronautical Conference* (pp. 1-5). Naples, Italy: International Astronautical Federation.

Hovenaars, A. G., Engelen, S., & Bouwmeester, J. (2012). Model-Based Discrete PID Controller for CubeSat Reaction Wheels Based on COTS Brushless DC Motors. *Proceedings of the 1st IAA Conference on Dynamics and Control of Space Systems* (pp. 379-394). Porto, Portugal: Univelt, Inc.

Klein, J. M., Budianu, A., Bentum, M. J., Engelen, S., & Verhoeven, C. J. (2013). Design of an Electric Power System with Incorporation of a Phased Array Antenna for OLFAR. *The 64th International Astronautical Congress (IAC 2013)*, (pp. 1-6). Beijing, China.

Quillien, K. A., Engelen, S., Gill, E. K., Smith, D. M., Arts, M. J., & Boonstra, A.-J. (2013). Astronomical Antenna for a Space Based Low Frequency Radio Telescope. *27th Annual AIAA/USU Conference on Small Satellites*, (pp. 1-7). Logan, Utah, US.

Rajan, R. T., Engelen, S., Bentum, M., & Verhoeven, C. (2011). Orbiting Low Frequency Array for Radio Astronomy. *IEEE Aerospace Conference*. Big Sky, MT, US. doi:10.1109/AERO.2011.5747222

Verhoeven, C., Monna, B., Engelen, S., Noroozi, A., Sundaramoorthy, P., Bentum, M., & Meijer, R. (2011). Emerging Eco-system: Nano-satellite Swarms and Large Satellites. *IAC 2011, IAC-11, D1,1,6,x11045*. Cape Town.

APPX. C: ABOUT THE AUTHOR

Steven Engelen was born in Hasselt, Belgium, on the 5th of April 1982. He attended secondary school from 1994 to 2000 at the Koninklijk Atheneum in Genk, Belgium.

From 2000 to 2009, he studied Aerospace Engineering at Delft University of Technology in the Netherlands. In this period, he has been active at DARE (Delft Aerospace Rocket Engineering), which is an amateur rocketry society, of which he is now a “Member of Merit”. During his time at DARE, he designed, built and launched several small rockets. Initially he was responsible for the design of the rocket body, as well as the control electronics. Later on his focus shifted towards more advanced liquid and hybrid bipropellant rocket engine development. This development resulted in four conference publications, and two journal publications.

He finished his M.Sc. degree with a thesis on radio-pulsar based navigation. From 2010 to 2014 he worked as a Ph.D. student at the chair of Space Systems Engineering at Delft University of Technology, and has been involved in the system’s engineering aspects of the OLFAR mission. In-between starting his Ph.D. research and finishing his M.Sc. degree, he designed several components deemed critical to attitude determination and control systems for nano-satellites. These designs he later on turned into products, initially under the umbrella of Syspa B.V., which was later spun out into Hyperion Technologies B.V., of which he is one of the founders, and currently functions as its CTO.

Satellite swarms are a novelty, yet promise to deliver unprecedented robustness and data-collection efficiency. They are so new in fact that even the definition of what a satellite swarm is is disputable, and consequently, the term "swarm" is used for practically any type of distributed space architecture.

This thesis poses the proposed definition of a satellite swarm as "a space system consisting of many **egalitarian** spacecraft, **cooperating** to achieve a **common global goal**".



Methods for designing such swarms are proposed and analysed, as well as the purported robustness and reliability commonly associated with swarms. The investigations show that, like with many systems, it is possible to create a swarm that is less reliable than even a single satellite, yet it is also possible to create one that is more reliable. However, this requires a paradigm shift, as in order to achieve this goal, a satellite swarm's satellites should be built as simple as possible, and this implies without internally redundant systems.

The OLFAR (Orbiting Low Frequency Antennas for Radio astronomy) mission, studying astronomical phenomena at low frequencies, has been used as a test case throughout the thesis, and various technological hurdles required for achieving the OLFAR mission are investigated and solved. This shows that while the OLFAR swarm itself is still slightly beyond current-day technologies, it is not as far out as originally thought, and it could well serve as a prime example of a mission for which a satellite swarm not only would be beneficial, but almost imperative.