

Fault-tolerant satellite computing with modern semiconductors Fuchs, C.M.

Citation

Fuchs, C. M. (2019, December 17). *Fault-tolerant satellite computing with modern semiconductors*. Retrieved from https://hdl.handle.net/1887/82454

Version:	Publisher's Version
License:	<u>Licence agreement concerning inclusion of doctoral thesis in the</u> <u>Institutional Repository of the University of Leiden</u>
Downloaded from:	https://hdl.handle.net/1887/82454

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <u>http://hdl.handle.net/1887/82454</u> holds various files of this Leiden University dissertation.

Author: Fuchs, C.M. Title: Fault-tolerant satellite computing with modern semiconductors Issue Date: 2019-12-17

Chapter 3

The Space Environment

Physical Fault Profile and Operational Considerations

A satellite's on-board computer has to cope with unique challenges which on the ground are only encountered in irradiated environments such in proximity of a nuclear reactor. Hence, for the understanding of the fault profile and application constraints for this thesis, in this chapter we provide an in-depth discussion of our operating environment and its effects on a satellite's on-board computer. We discuss the physical design restrictions aboard spacecraft, and operational considerations. Most importantly we discuss the impact of radiation on semiconductors, and how it can be mitigated.



3.1 The Impact of the Space Environment on Electronics

Space is a challenging environment for electronics to operate in, and its fault profile differs from that of most ground-based applications. A system engineer and computer architect has to consider many different design challenges and a very special fault profile. Only then is it possible to develop a reliable computerized system suitable for operation in this environment for an extended period of time.

3.1.1 Radiation Effects

Radiation is the main cause of faults in electronics aboard a satellite due to defects caused by electro-static discharge effects (ESD) [133] and directly inflicted particledamage. About 20% of all anomalies [134] aboard satellites can be attributed directly to high-energy particles, with the share of faults radiation-induced faults in electronicsheavy subsystems increasing drastically. This makes sense considering that the semiconductors used in electronics-heavy and computerized subsystems are more vulnerable to radiation-induced faults than, e.g., deployable structural elements (DE/STR) or a solar cell, whose performance will degrade slowly over time due to radiation.

A satellites on-board computer (OBC), communication transceivers (COM), and the electrical power system (EPS¹), as well as its attitude determination and control or orbit control system (ADCS/AOCS) all consist of microntroller- or processor SoCs with a varying set of peripheral electronics attached. A majority of all faults aboard CubeSats can be traced back to the failure of these architecturally similar subsystems [2] even directly after launch. Statistics from [2] are depicted in Figure 12. The low

 $^1\mathrm{Responsible}$ for battery charging and health control, as well as power management and distribution across a spacecraft.



Figure 12: Failures sources aboard CubeSat in 2016 after deployment, and after 30 and 90 days from [2]. Upon deployment, 61% of failures can be traced to strongly computerized subsystems. From a computer architecture perspective, all these subsystems are based on the same kind of components: non-fault-tolerant microcontrollers and mobile-market SoCs. After 90 days, 86% of all failures of CubeSat can be attributed to failures in the indicated subsystems. The base of data used by Langer et al. only serves as tentative indicator, as not all CubeSat and especially commercial operators choose to share this information.



Figure 13: A visualization of the three main natural sources of radiation affecting spacecrafts.

number of satellite failures due to payload (PL), or ADCS malfunctions in Figure 12 can be attributed to the fact that the failure of this subsystem seldom causes a satellite be lost entirely. Instead ADCS or AOCS failure will prevent certain mission objectives from being accomplished, the effects of which fail in the *early failure* and *partial mission success* categories as defined in Swartwout's CubeSat Database [42]. See also Figure 6.

Highly charged particles originate from a variety of different sources, which are depicted in Figure 13. They travel spinning around the Earth's magnetic field-lines in the Van Allen belts, are ejected by the Sun during Solar Particle Events, or arrive as Cosmic Rays from beyond our solar system. Galactic cosmic rays from beyond our solar system are mostly protons [5, 135], whereas various other high-energy particles are ejected by the Sun during solar particle events (*proton storm*). The radiation environment near the Earth, as well as in the rest of the solar system changes dynamically over time. We refer to this as *space weather*.

Depending on the orbit of the spacecraft and the occurrence of solar particle events, an OBC will be penetrated by a mixture of high-energy protons, electrons and heavy ions. In LEO, the residual atmosphere and Earth's magnetic field provide some protection from radiation, but this absorption effect diminishes quickly with altitude. Hence, microelectronics are exposed to a mix of highly charged particles, with flux density depending on solar activity and the spacecraft's attitude.

These particles can corrupt logical operations, induce bit-flips within data-storage cells (Single Event Upset – SEU) and connecting circuitry, or induce a latch-up. They can also cause displacement damage (DD), molecular changes in a chip substrate's crystalline structure which can cause its electrical properties to change, potentially causing permanent malfunctions. The particle flux will be increased while transiting the South Atlantic Anomaly (SAA), which is also depicted in Figure 13 [136]. Earth's magnetic field experiences a local, height-dependent dip within the SAA, due to the offset of the spin axis from the magnetic axis. In this region, a satellite and its electronics will experience an increase of proton flux of up to 10^4 times (energies > 30 MeV) [5]. This flux increase results in a rapid growth of bit errors and other upsets in a satellite's OBC.

Radiation challenges OBC fault coverage constantly throughout a mission and



Figure 14: The impact of radiation on a semiconductor varies depending on the used manufacturing technology. Manufacturing in fine technology nodes such as FinFET reduce the overall likelihood to experience radiation faults that affect critical logic due to shrinking geometry and therefore a reduced footprint of vulnerable logic. COTS Techniques such as FD-SoI furthermore increase SEE performance. Therefore a combination of small feature size manufacturing and robust COTS manufacturing in conjunction with software measures can offer strong fault tolerance capacity.

Image Credit: [138], Boeing/US-DTRA, for public use.

affects all of an OBC's components. The impact of radiation on different microfabrication processes, substrates, and memory technologies varies, as depicted in Figure 14. In general, electronics with a large feature size are more resilient to radiation-induced single event effects (SEEs) than those manufactured in finer production nodes. Chips with a small feature size are more susceptible to multi-bit upsets (MBU), that can propagate within circuits corrupting larger circuits or memory cells. Radiation events can also cause Single Event Functional Interrupts (SEFIs). These can affect sets of circuits, individual interfaces, or even entire chips. The cumulative effect of charge trapping in the oxide of electronic devices (total ionizing dose – TID) further impacts the lifetime of an OBC. Other types of radiation-induced faults, the destructive ones being the most relevant, are well described in [137].

As depicted in Figure 14, the robustness of a semiconductors in regards to different types of radiation-induced faults varies as well. Devices manufactured in old technology nodes with a coarse feature size show low TID (yellow line) and latch-up performance (blue), are robust to SEEs (green). Non-fault-tolerant semiconductors manufactured with old technology nodes are thus robust to SEE, while TID and latchup performance has to be increased through radiation hardening. CubeSat developers attempt to apply this same approach at the system level with modern semiconductors in the range well below the 50nm scale. However, SEE performance worsens with shrinking feature size, and drops below an acceptable level with modern technology nodes developed after the early 2000s. For comparison, commercial chip manufacturing using 130 nm technology nodes began in 2001, whereas at the time of writing smartphones-SoCs are manufactured with technology nodes between 16nm and 7nm. CubeSats seek to apply the latter kind of technology due to their much superior performance, lower cost, excellent availability, mature development tools, and reduced energy consumption.



Figure 15: Modern manufacturing techniques such as FD-SoI show much better performance under radiation than traditionally processes. Originally, these were developed to reduce feature size and energy consumption to achieve increase semiconductor packing density. Regarding radiation, the reduced footprint and inherent isolating properties of these technology nodes implies a reduced likelihood for a particle to induce an effect. With FD-SoI specifically, the changed structural properties of thereby manufactured chips further reduce the impact of SEEs due to the introduced an isolating layer of oxide.

Image Credit: Alles et al. [141].

In general, the effects of SEEs and SEFIs can be both transient and permanent, while DD is always permanent [5]. In case permanent effects are induced, or faults occur in memory, radiation induced faults accumulate over time. The accumulative nature of permanent faults implies accelerated and often spontaneous ageing, which must be handled efficiently throughout the entire mission.

The increased impact of SEE on finer feature size chips also invalidates the naive approach of achieving better protection by adding more circuit-level protection. This prevents the continued application of traditional RHBD/RHBM concepts [104,132] to modern, high-performance embedded and mobile-market SoCs. The energy threshold above which SEEs induce transient faults in chips manufactured in fine technology nodes decreases, and the ratio of events inducing multi-bit upsets (MBU) or permanent faults increases.

Radiation tests with FinFET [139] and Fully Depleted Silicon On Insulator (FD-SoI) [140] based technology nodes also show improved SEE performance, contrary to projections based on technology scaling. As depicted in Figure 15, transistors in these technology nodes have a much reduced footprint as compared to bulk manufacturing. The smaller feature size there reduces the likelihood for a charged particle to interact with sensitive chip regions, which results in fewer but more severe upsets in such semiconductors [141]. FD-SoI introduces an additional layer of isolating oxide, which helps reduce reduce the impact of radiation effects on such a semiconductor. Hence, chips manufactured in certain new technology nodes, such as recent generation FPGAs [142] show better than expected TID [143] and latch-up performance [144], while also showing different SEE performance: fewer non-masked events with more severe impact.

In practice, radiation induced faults may corrupt computations of a computer, corrupt register contents, data stored in caches, main memory, and non-volatile memory. Memory mainly suffers from bit-rot and malfunctions in controller logic, and for volatile memory, these can well be compensated for using error correcting codes (ECC) combined with error scrubbing. Non-volatile memory also requires more powerful era-



Figure 16: The functional principles and structure of two of the currently most promising inherently radiation immune memories: PCM (a), and MRAM. Radiation immunity of these cells is based on that these memories do not store information as a charge, in contrast to radiation-susceptible DRAM, SRAM, or Flash.

Image Credit: (a) Hayat et al. [145] (b) Fert et al. [146].

sure coding systems, the basic notions of which also exist in latest-generation COTS flash memory based devices for ground use, as there galactic cosmic rays have become relevant sources for faults due to technology node scaling.

Functional interrupts can cause individual processor cores or other sub-units of a semiconductor to fail temporarily or permanently. Data can also be corrupted in transit, e.g. while being transferred or due to upsets in peripheral interface controllers. Hence, from a developer's perspective, to-be executed software and data can only be considered fault-free if it resides exclusively in radiation-hard memory and radiationhard processing logic throughout. As this is not the case with all but trivial processing logic, no part of an OS can be relied upon to be fault-free, and concepts requiring such an entity do not offer effective fault coverage in the space environment.

The memory cells of certain novel memory technologies (e.g., MRAM [147], and ReRAM [148], and PCM [149]) have been shown to be inherently immune. This is due to the data storage mechanism in these non-charge based memories [150, 151]. The memory cells of commercial MRAM and PCM ICs are largely immune to radiationinduced faults, and their structure and operating principle is depicted in Figure 16. However, connecting circuitry and controller logic of these parts if still vulnerable to radiation, and incorrect addressed memory can very well cause data corruption during read and write operations.

Flash memory, one of the most widely used charge-based memory technologies, has been shown to be rather susceptible to radiation effects [153]. Each flash memory cell contains a single field effect transistor with an additional floating gate, which is depicted in in Figure 17. Voltage applied between source and drain generates an electric field with a conductive channel through which electrons can flow into the floating gate. The state of a cell is thus dependent on whether or not a specific threshold voltage is exceeded (programmed, Figure 17b) or not (erased, Figure 17a).

Radiation can induce a variety of effects in charge-based memory such as flash [153]. In Figures 17c and 17d, we depict two opposing effects induced by particles with a positive and negative charge [154]. In Figure 17c a cell in erased state is hit by a

negatively charged particle. Such a particle can cause a storage cell to change its state by depositing electrons in the floating gate as it passes through the structure. Figure 17d depicts the inverse effect with a positively charged particle, which changes the net charge of the floating gate. The particle event may cause the charge in the floating gate to rise or drop one rise above or drop below a volatile threshold of the cell and thereby change the value represented by the storage cell.

Particles may also alter the structural integrity of different parts of the memory cell, e.g., draining the gate, or causing permanent damage [153]. Due to a shifting voltage threshold in floating gate cells caused by the total ionizing dose, flash memories become more susceptible to data degradation due to leakage. Modern multi-level cell flash memories manufactured in fine technology nodes are more prone to SEUs causing shifts in the threshold voltage profile of one or more storage cells [153]. Flash cells can also store more than a bit of data, and then also become susceptible to



(c) Erased cell hit by a negatively charged particle (d) Cell reset by a positively charged particle

Figure 17: The structure of a Flash memory cell in erased (a) and programmed state (b), inspired by a figure from Zandwijk et al. [152]. Data is stored as charge in a floating gate attached to a controlling field effect transistor. Radiation can induce a variety of different effects in charge-based memory [153], and in Figures (c) and (d) we depict two opposing effects induced by particles with a positive and negative charge [154].

MBUs: radiation may cause a state change across multiple voltage levels [155]. The semiconductor's temperature and particle events can also influence the leakage current of a these memory cells, thereby reducing the charge stored within the floating gate over time [156]. The radiation-induced effects depicted in Figure 17 are representative for the entire class of charge-based memories, even though other memory technologies store data as charge in electrically different ways [157].

Physical shielding using aluminum and other materials can reduce certain radiation effects [158]. The necessary shielding strength depends on the physical properties of the material used for shielding [159]. This approach has been used extensively in classical space applications in the early time of spaceflight. However, the level of shielding needed to protect modern semiconductors from radiation effects would require a miniaturized spacecraft to dedicate an unreasonable additional mass and volume to shielding [159]. For very large satellites, the use of strong shielding is still a viable (but costly and inefficient) option [160].

Weak shielding can introduce scattering effects, while offering nearly no added protection [161]. These can occur due to interaction of a highly charged particle with shielding material, which can cause a shower of charged secondary particles. This secondary particle radiation takes the shape of a cone from between the point of impact of the original particle and the underlying semiconductor [161]. Particle scattering can therefore cause multiple particles with lower charge to penetrate a semiconductor, instead of just one. Hence, very thin shielding such as aluminium-RF-cages commonly found in consumer electronics offer usually no radiation protection [159].

3.1.2 Design Constraints for Space Electronics

The success of a satellite missions depends on designer's ability to develop a system that can withstand operation in the space environment, and can cope with the design constraints that are in place aboard a satellite. In the remainder of this section, we therefore provide a brief overview of satellite design constraints.

Solar cells are the main power source aboard modern spacecraft in the inner regions of the solar system [7]. A spacecraft's orbit, location and orientation (attitude) relative to the Sun, and the solar array's temperature all influence the efficiency of its solar array. Miniaturized satellite's have small solar arrays with varying output, and their OBCs are limited to a few Watts of power-budget (power consumption averaged over time).

Operation in the space environment outside planetary atmospheres means that a satellite will operate in vacuum [162]. In turn, this implies the absence of the heat-transfer medium necessary for thermal convection, and hence also air cooling. Depending on the specific chip design implemented within a semiconductor, this can cause a chip and its packages to exhibit different or even anomalous thermal properties, potentially causing hot-spots and impact performance and lifetime [163]. Heat generated within a spacecraft therefore has to be transferred to the exterior and is then emitted as infrared radiation. A variety of engineering measures are available to help create a stable spacecraft-internal temperature environment [164].

Operation in vacuum and the low temperatures encountered in the space environment, can cause rapid material aging. The extreme temperature deltas when operating in a planetary orbit in direct sunlight and darkness can furthermore cause out-gassing, e.g., of chemical softeners present in materials such as plastics [165]. Gassed-out chemicals may interact with other components of a spacecraft, especially sensors, and may cause folded solar cell arrays to stick together, fold incorrectly, and fail to deploy from stowage [166]. This effect is a major problem for spacecraft equipped with optical payloads, e.g. astronomical observatories: out-gassed chemicals may then accumulate over time on sensors, mirrors, and lenses, and degrade an instruments performance. In large spacecraft projects, components are therefore often baked at high temperatures or exposed hot-cold cycles to reduce this effect in space as much as possible.

Upon launch, satellites have to withstand considerable physical stress and may experience vibration-induced resonance effects [167]. To a certain extent, these can be simulated through mechanical means (shakers) and acoustics on the ground, and then mitigated through engineering and a wise choice of materials. To design computer systems to better cope with launch stress and the extreme temperature changes that may be encountered in the space environment, electronics can be packaged in more suitable materials than the usual plastic packages used on the ground. However, electronics in ceramics and metal-based packages are at the time of writing significantly more expensive than conventional consumer parts, and usually non-options for CubeSat applications. Specialized materials can also be used in the different layers of a PCB, and can help optimize electrical, structural, and thermal properties, which today is also used aboard miniaturized spacecraft, e.g., aboard the MOVE-II CubeSat.

3.2 Technology Readiness and Standardization

Satellite missions can last from several months up to many decades, and therefore satellite designers may encounter hard technological barriers such as data retention [168]. Examples include, but are not limited to, issues with using electronics storage technologies due to limited data retention periods, solar cell degradation, and material degradation due to long-term thermal stress and out-gassing.

Traditional space companies and organizations are very cautious when considering new technology with little or no space heritage. Often, they modify and adapt existing, foreign industry standards to their own needs instead of reusing them, and develop their own standards [169]. Several sets of space related quality and design standards exist, which are administered by committees consisting of space agencies, governmental bodies, military and major industrial actors. Some of these standard libraries are published, while others remain proprietary (e.g., ARINC) or are even kept confidential (military standards). Currently, the most relevant publicly available and widely adopted standards are published by the *Consultative Committee for Space Data Systems* (CCSDS), the *European Cooperation on Space Standardization* (ECSS), and the *NASA Technical Standards Program*. Standards popular in the IT-industry in general do influence avionics design (e.g., Ethernet/IEEE 802.3 is today the technological base for AFDX [94], but adoption of this technology has taken more than 30 years), but mostly indirectly due to a technological lag between IT-industry and space-avionics that ranges from between 10 to 40 years [170].

Avionics (thus, Aerospace and Spaceflight electronics) development relies not just upon specialized and tested components. Instead, technological maturity has to be proven in practice to demonstrate that a component or technology is ready for application in the space environment. Thereby, the quality and heritage of a solution are assessed based on a standardized set of indicators resulting in a classification in technological readiness levels (TRLs) [171], see Figure 18. For some types of chips the



Figure 18: Technology readiness levels and the requirements to qualify a component for a certain level.

Image Credit: NASA, public domain.

global space industry may have annual demands for only several hundred or thousands of chips, resulting in extreme per-device development costs compared to common IT industry production quantities of millions of units. Due to limited alternatives and their requirement to rely upon proven and validated hardware, the space industry and their customers must afford high hardware costs and accept long development cycles [169]. The TRL required for a component may vary per usage scenario and subsystem, the highest level is thus not necessarily required and TRL9 components may even be replaced with less expensive or maybe more modern components with a lower level.

In this thesis, we propose an architecture which incorporates a set of theoretical fault tolerance concepts, which exist at TRL1. Based on these concepts, we formulate a conclusive architecture for our application, which initially exists at TRL2. We then proceed to conduct fault-injection experiments with a proof-of-concept of the architecture (TRL3), and produce a practical implementation based on development-board components in a bread-board setup (TRL4).

3.3 Operational Constraints for Satellite Computers

In contrast to most earth-bound computing, it is not possible to physically access a spacecraft in orbit [172] to diagnose or resolve faults. However, this does not mean that they can not be repaired, refueled, upgraded, or otherwise serviced during a mission. In fact, most spacecraft are designed to be service friendly, as this makes them easier to assemble and test on the ground. This is especially important as testing of a spacecraft as a whole and its individual subsystems is a complex and costly undertaking. Component-level as well as testing of a full avionics system makes up a significant share of the time needed for the design and construction process.

Hands-on maintenance or diagnostics on-orbit are uncommon today, and servicing missions have been conducted only on a few occasions. All of these spacecraft were large satellites and space-stations in LEO with outstanding significance to science, society, or driven by national interests. Prominent examples include the Hubble Space Telescope [173] and several space stations [172,174,175], where servicing was required to resolve faults. For most modern non-agency and non-governmental satellites, and especially smaller and cheaper spacecraft, hands on maintenance is not feasible, and usually also not economical [173]. Hence, an on-board computer has to operate and handle faults autonomously over the entire duration of a spacecraft's mission, which may last for several decades.

Diagnostics of computerized systems therefore have to be conducted remotely and in a scripted manner locally aboard a satellite. Considering the journey of Cassini/Huygens depicted in Figure 19, this implies differences in link behavior and communication bandwidth during a mission. Even in earth orbit, a satellite's telemetry and telecommand (TMTC) link is lossy, and offers very low bandwidth compared to ground-based communication (in the low kbps range). As depicted in Table 2, signal travel times in LEO and Geostationary Earth Orbit (GEO) still allow widely used network communication protocols for ground use to be utilized, if aspects such as Doppler-Shift are compensated for [178].

All CubeSats launched until 2018 operated in a LEO [17], and most utilize a combination of UHF and VHF frequency bands to realize their commandeering channel. LEO communication windows between a ground station and a satellite are limited to between 5 and 20 minutes in ideal weather conditions, and reduced by equipment dampening, environmental effects, and atmospheric conditions [179]. Only part of this communication window allows actual communication with a spacecraft due to link-quality issues. The actual duration varies depending on the satellite's orbit and the environment the ground station operates in: buildings, natural obstacles and fading signal quality with declining elevation angle when approaching the horizon all affect a link's signal-to-noise ratio [180]. For comparison, while commandeering the FirstMOVE CubeSat, actual link availability during communication windows never exceeded 12 minutes.

LEO-link availability can be increased through the use of satellite-relays (e.g., TDRS [181]) and ground-station networks [182]. However, these are currently largely unavailable to miniaturized satellites due to economical considerations on the operator's side and form-factor and cost constraints for miniaturized satellites. In practice, this curtails remote debugging capabilities of spacecraft. It prevents the direct re-use of, e.g., all low-level testing protocols which are today widely used on the ground application such as JTAG or ICE, and prevents remote-debugging using standard debugging





tools.

When communicating with spacecraft orbiting other planets in our solar system, signal travel times and thus link latency grow rapidly. With space probes traveling beyond the Earth/Moon system, the available link rates decrease sharply and often only few hundred bps can be achieved. Unidirectional signal travel times to neighboring planets make real-time bi-direction communication concepts as used on the Earth technically impossible. At the time of writing, the mars rover Curiosity can achieve a data rate of between 500 bps up to a theoretical maximum of 32 kbps and round-trip times of at least 8 minutes under ideal circumstances [183]. The TMTC link of the Voyager probes [184] can achieve a maximum of 160 bps at the edge of the solar system via the Deep-Space Network [185] with signal travel times approaching a duration of a day.

As depicted in Figure 19, a spacecraft may have to travel within our solar system for years, before actually arriving at its destination, where it can then begin to perform its actual mission. During such missions, the performance requirements to a satellite computer can vary. In Figure 20, we depict a simplified version of the orbit/work schedule of NASA's Enceladus Life Finder (ELF) probe, which will conduct science on Saturn's sixth largest moon. Travel to the Saturn system will take years, but once

Communication	munication Distance from Earth		Signal Travel Time	
Endpoint	Min.	Max.	Min.	Max.
LEO	$400 \mathrm{~km}$	$2,000 \mathrm{~km}$	$3 \mathrm{ms}$	18 ms
GEO	$35{,}786~{\rm km}$	-	${\sim}250~{\rm ms}$	-
Moon	$356{,}400~\mathrm{km}$	$406{,}700~{\rm km}$	$2.4 \mathrm{~s}$	2.7 s
Mercury	$0.62 \ \mathrm{AU}$	1.39 AU	$5 \min$	$12 \min$
Venus	$0.28 \ \mathrm{AU}$	$1.72 \ \mathrm{AU}$	$2 \min$	14 min
Mars	$0.53~\mathrm{AU}$	$2.52 \ \mathrm{AU}$	$4 \min$	$21 \min$
Jupiter	$4.21 \ \mathrm{AU}$	$6.21 \ \mathrm{AU}$	$35 \mathrm{~min}$	$52 \min$
Saturn	$8.54~\mathrm{AU}$	$10.54 \ \mathrm{AU}$	1:11 h	1:28 h
Uranus	$18.23 \mathrm{AU}$	$20.23 \mathrm{AU}$	2:32 h	2:48 h
Neptune	$29.06 { m AU}$	31.06 AU	4:02 h	4:18 h
Voyager 2	$\sim 121 \text{ AU}$	-	${\sim}16{:}50~{\rm h}$	-
Voyager 1	${\sim}147~{\rm AU}$	-	${\sim}20{:}22~{\rm h}$	-

Table 2: Unidirectional signal travel times for radio communication in vacuum between a ground station and a spacecraft at a particular location in the solar system. Distances between the Earth and different planets in the solar system vary due to celestial mechanics. In practice, the signal latency even for LEO communication is drastically larger than the theoretical signal travel speeds indicated here due to latency in the signal processing chain. Data for the Voyager probes based on https://voyager.jpl.nasa.gov/mission/status, accurate as of September 2019.



spacecraft's on-board computer. properties. In each color-highlighted orbit-segment, ELF has to conduct a different task or operation, with different requirements towards the Figure 20: NASA's Enceladus Life Finder (ELF) is scheduled to make ten flybys of Saturn's moon Enceladus to investigate that its environmental

Image Credit: Figure from public-use preview press material by the ELF Team/JPL/NASA, Based on a figure from [186]

ELF has entered orbit around Enceladus, it will have to handle a variety of different tasks with very different system requirements (indicated in color). We utilize this satellite's mission operations schedule to highlight how requirements to a satellite's on-board computer can shift during a mission.

During the yellow-outlined communication phases, reliability of the satellite computer is crucial, as communication windows are brief and the available link-rate is low. Any lost communication time could directly impact the satellite's mission and subsequently executed tasks. Ideally, during this time a satellite's computer should offer increased fault tolerance capabilities at the expense of other system parameters, if such capabilities were available.

The red- and purple highlighted orbit segments indicate times when ELF will perform maneuvers through its propulsion subsystem and adjust the orientation of it's solar panel array. When performing maneuvers, precise timing and therefore the ability for real-time operations are crucial, while overall compute performance requirements will be comparably low. Finally, in the green-market science phase, performance is critical, and during this phase, spending extra energy to increase the satellite's overall compute and data-storage capacity may allow the spacecraft to conduct more and better science within its brief mission. With the computer architectures used aboard spacecraft today, little adaptivity is possible. However, future satellite computers based on modern mobile-market and embedded computer architectures could very well support such functionality if fault tolerance capabilities can be adjusted at runtime.