FY23 Strategic Initiatives Research and Technology Development (SRTD)

High-Performance Fault-Tolerant Compute Element for Planetary Science Missions

Principal Investigator: Marcel Schoppers (348); **Co-Investigators:** David Foor (319), Dennis Kou (348), Diego Martinez (174), Gregory Miles (332), Daniel Nakamura (349)

Strategic Focus Area: High-performance fault-tolerant Compute Element for next decadal planetary science missions | **Strategic Initiative Leader:** Jeffery L Hall

Objective: A computing architecture that delivers both the high reliability required for inter-planetary flagship spacecraft, and the latest-COTS computing performance, *cheaply upgradeable*. This is being achieved by using space-grade computers to manage multiple "redundant" high-performance but low-cost Systems-on-Chip as co-processors, of which several may fail, without harming the mission. Our specific objectives are

1. easily upgraded with commercial state-of-the-art compute elements,
2. high reliability and availability, suitable for planetary EDL and helicopter flight,
3. fail-operational despite 2 co-processor failures or data-network failures,
4. computing throughput 1000× a traditional RAD750 processor,
5. a small SWaP footprint (e.g. < 30W) suitable for the SmallSat class of missions,
6. an interface board matured to TRL-5 (see Technology Readiness Levels),
7. redundancy- and fault-management software matured to TRL-5.

Background: The Planetary Science Directorate is currently formulating the next generation of missions, following Europa Clipper and Mars Sample Return, to explore our solar system with e.g. the Endurance-A lunar rover, a Mars Science Helicopter (at right), a Venus Aerobot, and many missions to smaller bodies. These future missions will require significantly higher computing power, for Terrain Relative Navigation (during EDL and helicopter flight), auto-nav, and onboard planning, scheduling, and science-data processing. Furthermore, we would like to become ready to exploit computing capabilities that will certainly continue to improve.

Approach and Results: The proposed architecture consists of 1–4 COTS high-performance but non-rad-hard compute elements, being managed by 1–2 highly-reliable rad-hard compute elements, with a MiddleWare layer to recover the needed fault-tolerance, reliability, and availability. The compute elements are interconnected via a key Interface Board we are building, which is passive, and theoretically agnostic to the choice of compute elements, to allow us to henceforth exploit the commercial state-of-the-art in computer hardware.

A Triple Modular Redundant configuration operates-through a permanent fault occurring anywhere in the system. A Quadruple Modular Redundant configuration operates-through failure of two high-performance elements. If the failures are sufficiently separated in time, the MiddleWare supports recovery to full health (this currently requires software-engineering for each task/thread being recovered).

The key tasks for FY2023 were, designing and fabricating the Interface Board to accommodate compute elements we have in-house, and maturing the MiddleWare.

Significance/Benefits to JPL and NASA: The proposed architecture is expected to provide flagship-class C&DH/GN&C reliability, plus onboard science data and autonomy processing throughput, plus fault tolerance against all 1-fault and some 2-fault scenarios. It will also allow un-interrupted operation during mission-critical phases, for significantly lower SWaP and cost than the state-of-the-practice in spacecraft computers.

In the near-future, this technology can benefit the Mars Science Helicopter and the Endurance-A Lunar Rover. It will also enable # flagship spacecraft with state-of-the-art computing, # highly-reliable smaller spacecraft, # improved reliability in radiatively hot environments, # precision landing with hazard avoidance done by COTS co-processors, # rover auto-nav done by COTS co-processors, and # aerial flight on Venus, Mars, the outer planets, and Titan.

The notional Mars Science Helicopter, during EDL, ejected from the aeroshell in mid-air, to land itself. The computers must tolerate potential radiation hits, find a safe place to land using machine vision, and deliver > 99.99% availability in flight.
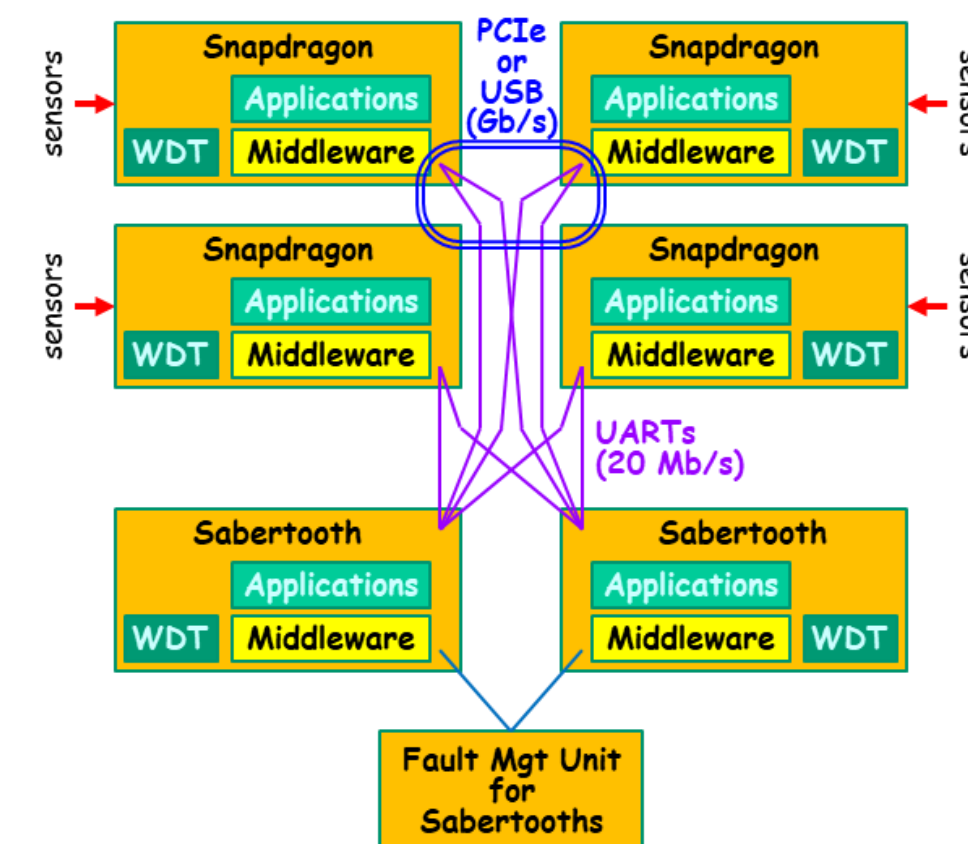
The multi-compute element architecture, showing 4 high-performance elements as Snapdragons, 2 highly-reliable elements as Sabertooths, and MiddleWare running on each. The Interface Board adds the three networks (blue, purple, green), power switching, and fault containment.
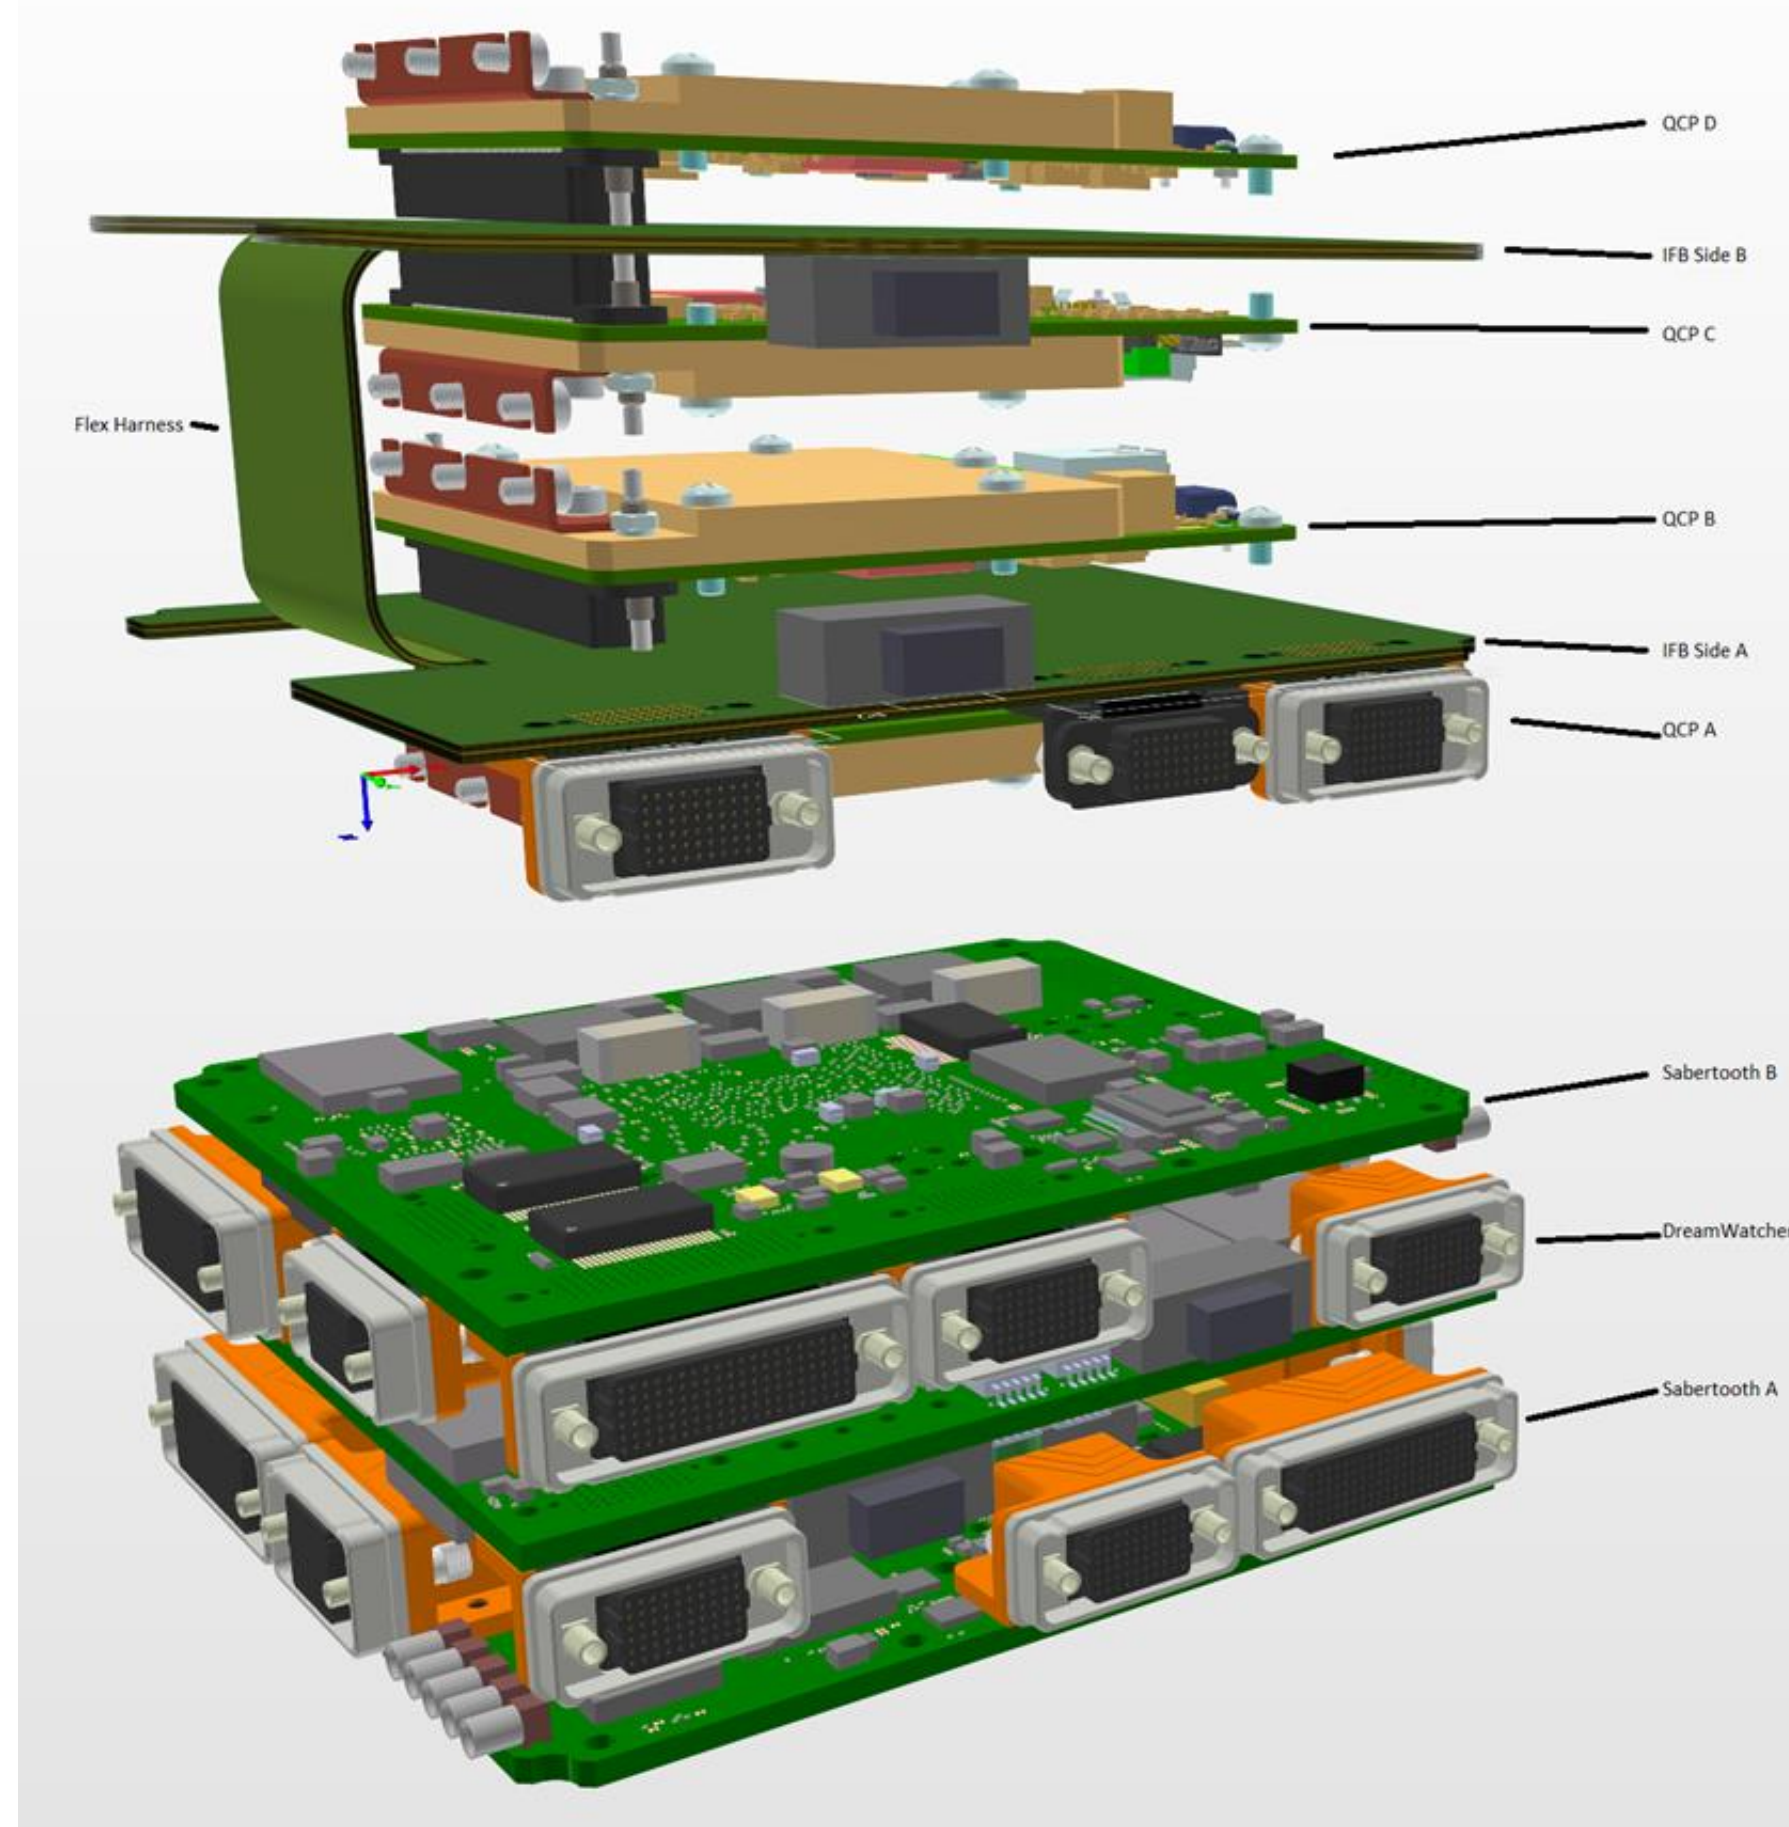
Design of the multi-computer, approximately 1U in size. Top: stack-up of 4 Snapdragon CoProcessor boards (tan), and the Interface Board in 2 halves (green). Bottom: stack-up of 2 Sabertooths and their Fault Management Unit in the middle (green). Not shown: cables for power, and for UARTs between Sabertooths and Interface Board.

PI/Task Mgr. Contact Information:
818-354-9290 marcel@jpl.nasa.gov