

MarsIPAN: Optimization and Negotiations in Mars Sample Return Scheduling Coordination

Jasmine T. Otto , University of California, Santa Cruz, CA, 95064, USA

Malika Khurana , The New York Times, New York, NY, 10018, USA

Noah Deutsch , Harvard University, Cambridge, MA, 02138, USA

Benjamin P. S. Donitz , California Institute of Technology, CA, 91109, USA

Oskar Elek , University of California, Santa Cruz, CA, 95064, USA

Scott Davidoff , California Institute of Technology, CA, 91109, USA

Resource allocation problems touch almost every aspect of modernity. We examine communication bandwidth optimization and negotiation in NASA’s early stage Mars Sample Return (MSR) mission, which places multiple robots into a single region on Mars. We present a design study conducted over two years at the NASA Jet Propulsion Laboratory with MSR, which characterizes the design and evaluation of the deployed MarsIPAN schedule browser. We find that MarsIPAN changes the schedule analysis process, providing new insight into how bandwidth is allocated and enabling individual spacecraft teams to openly negotiate for scarce resources. This visibility leads to changes in how spacecraft teams apportion bandwidth, plan staffing, and organize and share resources. We reflect on the design study methodology as revealing, documenting, and supporting existing communication processes and software infrastructure within knowledge-intensive organizations.

Scheduling problems are common in domains ranging from manufacturing¹ to logistics and operations research.² Consequently, scheduling optimization offers a rich laboratory to explore sociotechnical negotiation over finite, economized, and high-demand resources mediated by complex technical systems. This article presents a visualization design study, which explores optimal scheduling in the very early stage Mars Sample Return (MSR) mission.³ In order to support ongoing discussions around the extent and feasibility of planned MSR operations, we developed the MarsIPAN schedule visualization system, which enables stakeholders to explore optimal pass bandwidth allocations.

MSR will put extreme stress on the deep-space communications infrastructure by situating multiple spacecraft, each with its own independently varying bandwidth requirements, into a single region on the Mars surface. Because each spacecraft depends on the same scarce bandwidth to (and from) Mars, space agencies plan out bandwidth allocations years into the future to ensure successful mission operations.⁴ This represents a ripe area for visualization to support communication around scheduling optimization, and how operational complexity is “traded off” between various scientific and engineering stakeholders.⁵ Stakeholders use MarsIPAN to assess the quality of individual allocations and to explore the implications of different threat scenarios and contingency plans.

Over two years, the authors of this article partnered with aerospace, communications, and systems engineers at the NASA Jet Propulsion Laboratory (NASA/JPL) to investigate the impact of this visual analytics system on relay bandwidth and staffing

0272-1716 © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies.
 Digital Object Identifier 10.1109/MCG.2025.3558426
 Date of publication 7 April 2025; date of current version 8 August 2025.

discussions. In this article, we document the three visual-interactive schedule components comprising MarsIPAN. Together, these represent *operational cycles* formed by the relationship between the execution window on Mars, relay pass availability, and surface teams' bandwidth requirements on any given day. We characterize MSR experts' needs and tasks, illustrate two use case scenarios, and explore a single communications scenario in-depth with both a relay lead and a surface lead, surfacing their different concerns and approaches.

Scheduling problems are critical opportunities for visualization practitioners to enable informed interdisciplinary negotiations. We argue that shared visualization systems enable participants with different forms of expertise to exploit shared abstractions and to ground their respective approaches to common goals in diverse sources of data. Like visualization sandcastles,⁶ which add grounded theory analysis layers⁷ to visualization research artifacts, we encapsulate the complexity of our stakeholder collaboration into a *knowledge transit diagram*, representing theoretically how our multidisciplinary stakeholder groups create common ground, encapsulate and exchange expertise, and collaboratively elaborate a shared understanding as part of the overlap between their shared and competing goals.⁸ We see the successful visualization design study as a type of workplace ethnography, seeking to identify the role of each stakeholder team. Because these roles are connected through shared languages and information systems, despite differing in their disciplinaryity, the visualization system reifies shared terms for interdisciplinary collaboration.

This article makes the following contributions.

- 1) A two-year design study⁹ of collaborative schedule analysis on the MSR project, which requires three robotics teams to share relay communications bandwidth during coordinated Mars surface operations. This study characterizes the MarsIPAN schedule (the *Mars Interactive Pass Allocation Navigator*), a visual analytics system that enables users to reason about pass allocation scenarios.
- 2) A critical reflection on knowledge work underpinning MSR operations analysis. We embedded with a NASA/JPL team investigating communication needs across multiple teams around conflicting bandwidth requirements. By observing expert usage of the MarsIPAN schedule analyzer, we saw that stakeholder relationships within MSR correspond with and

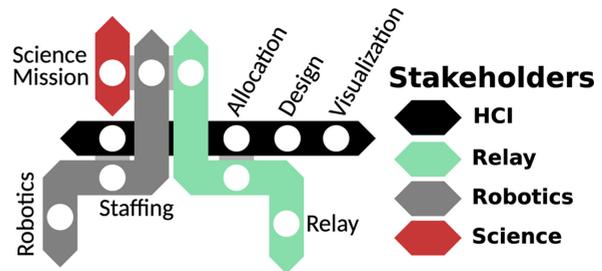


FIGURE 1. A knowledge transit diagram representing MSR communications at NASA/JPL. Each “line” corresponds to a domain of expertise, containing multiple “stations” (white dots) where individuals or teams conduct work. In MSR, relay operations (teal) and surface robotics (gray) domains both interact directly with scientists (red), who establish the mission requirements. These requirements are conveyed indirectly to HCI (black), which helps relay and robotics to reason in joint about communications feasibility. (Figure reproduced from Otto and Davidoff.⁸)

cut across the implicit disciplinary structure of NASA/JPL's vast data landscape.¹⁰ We document these relationships through a knowledge transit diagram (see Figure 1).

BACKGROUND

Scheduling and Optimization

Schedules are sequences of events used to coordinate actions of finite resources over a period of time. Schedule optimization problems are found across numerous domains connected to manufacturing¹ and logistics.² Communications bandwidth allocation is a special case of operations scheduling across multi-agent systems, which has been studied by NASA for decades.¹¹

The early stage MSR Mission³ introduces the challenge of supporting multiple spacecraft on the surface of Mars, each of which is operated by different teams with shared scientific but competing engineering goals. Different teams will work at different cadences, with differently complex spacecraft, each day requiring negotiation around the scarce allocation of uplink bandwidth (data going to Mars) and downlink bandwidth (data coming from Mars).

In addition to the scheduling problem itself, resource allocation problems across multiple teams are inherently wicked; everyone can find a use for that extra bandwidth. Thus, optimization possesses an inherent political dimension. Creating a joint

efficiency metric across multiple teams requires making explicit decisions about which team to prioritize; taking the average of three concerns is not the same as hearing three concerns. Visualization practitioners, then, are in an unusual position where both the metric representing efficiency and the underlying concern bear equal weight in their analysis.⁷

Scheduling and Visual Analytics

Schedules are a visual representation of event sequences along a timeline. Unpacking this definition reveals a surprising number of design decisions regarding time, events, and coordination. Between different applications, temporal resolution can vary greatly;¹² during active operations, surface rover activities are typically planned down to the minute on a daily basis.¹³ By contrast, each orbital relay pass takes about an hour to supply its entire bandwidth. These passes form a pattern that cycles over several months and is managed years in advance.

Different stakeholders take varying degrees of initiative when reading and interpreting an “optimal” schedule. An academic conference scheduling system was designed to let conference organizers reorder its proposed solutions based on crowdsourced knowledge of attendee interests.¹⁴ In aircraft line maintenance, managers applied their domain expertise around existing best practices and social mores to the staffing schedule, rejecting several optimized schedules to avert a breakdown in legal compliance.¹⁵

Although operations scheduling builds on deterministic problems, such as job-shop scheduling, because a few days of activity may include thousands of constraints, operators must also adapt the schedule to unpredictable events, such as scientific opportunities and machine breakdowns.¹⁶ In communications bandwidth planning, team leads take a high degree of initiative, using scheduling algorithms to support their own sense-making and pointing out oversimplifications present in the scheduling algorithm.

Lastly, visualization designers who create schedules do not just surface time- and event-ordered data, but also mediate and facilitate coordination around these representations so that teams can resolve potential resource constraints and/or conflicts within the mission. Visual analytics includes a design dimension that is sensitive to and ordered around organization dynamics, and often requires bespoke software¹⁷ to broker technical decision-making. This article contributes to that discussion, adding a transit diagram,

which codifies this negotiation dynamic, illustrating and building on earlier work⁸ on both how teams interact via mediated visual tools and how visualization designers encapsulate expertise to broker those transactions.

DESIGN STUDY

This design study⁹ investigates visual analytics for scheduling optimization in the context of the underdevelopment MSR mission.³ The current MSR mission architecture includes multiple spacecraft: a Sample Retrieval Lander (SRL), multiple Sample Retrieval Helicopters, a Mars Ascent Vehicle, and an Earth Return Orbiter (ERO). This architecture radically increases the local population of both surface and orbital robots. Because each robot will rely on solar power, each one needs to move around every day, so that accumulated dust doesn’t kill the robot by preventing it from charging; therefore, stakeholders must understand whether there is enough bandwidth to command each robot on each Mars day.

Teams at NASA/JPL, a large research laboratory participating in MSR, are engaged in ongoing bandwidth scheduling discussions, which will influence MSR mission design. Our visualization group was contacted by a team, which develops optimal bandwidth allocations to support ongoing and future negotiations between relay experts and surface operations teams.

For the first 10 weeks, our team of three visualization designers conducted formative user fieldwork with the *allocation design team*, who not only gave us access to their allocation data, but also introduced us to their working practices and the communication expectations of other aerospace experts.¹⁸ During this initial engagement, team visualization design team also iteratively developed early versions of the MarsIPAN path allocation schedule visualizer.

Over the next 18 months, our team met with the allocation design team over 60 times. During each 1-h session, they interacted with new versions of MarsIPAN for about 10 min (or more, if we had added more complexity to the interface than usual) and talked with us about feature improvements and other stakeholders’ needs for the rest of the time. This study-design-build-learn cycle allowed us to iteratively identify process, data, visualization, and interaction improvements to the visualization system itself. In addition, we mapped out other JPL stakeholders and their needs, which upon reflection, we would develop into the knowledge transit diagram (see Figure 1).

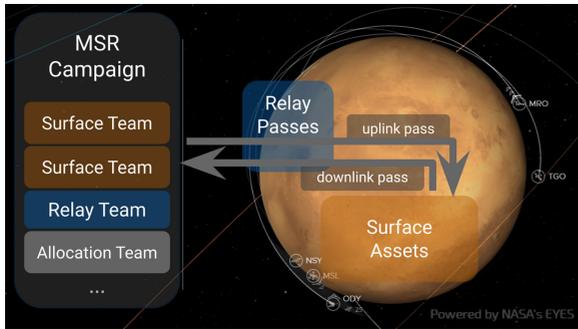


FIGURE 2. Conceptual map of MSR Campaign stakeholders and their respective operational responsibilities. The relay team (blue) provides pass bandwidth to and from Mars, which the surface teams (orange) need to operate their respective assets. The systems engineering team (light gray) enumerates the relay passes and recommends which ones should be used for uplink or downlink by each surface team, according to their needs.

In addition to working closely with the allocation team, we engaged in 10-h long co-design sessions with *relay and surface experts* over the 18 months, learning about their tasks and needs through their interactions with MarsIPAN. Finally, we conducted semistructured, end-user evaluations of MarsIPAN prototypes after six months and 12 months of development. In evaluation, we present a thematic analysis of the two interviews collected after 12 months, reflecting the mature MarsIPAN system.

Technical Requirements

The MarsIPAN schedule browser serves operations experts engaged in collaborative analyses supporting MSR campaign goals (see Figure 2). MarsIPAN is an Earthtime–Marstime schedule optimization visual analytics system driven by the previously described MARMLADE multispacecraft scheduling algorithm,¹⁹ which efficiently allocates *relay pass bandwidth* as a function of orbital relay availability under a set of possible operational constraints. The scheduling bandwidth team recruited us to improve their existing visualization system, which was manually drawn in Excel on a pass-by-pass basis.

MARMLADE allocations are parameterized by the Marstime window during which surface assets can execute instructions (because of energy and temperature considerations), and the Earthtime window during which the operations room is staffed (because of human factors constraints); together with the schedule of flyover passes provided by the

TABLE 1. Simplified version of the data schema for each relay pass allocation. The Orbital Asset and Surface Asset are typically specific vehicles; one in orbit, and one on Mars. The Sol and LMST (Mars time) indicate when the pass occurs, while the assignment indicates whether it is used for uplink or downlink bandwidth. Finally, the FWD and RET volumes indicate how much bandwidth is provided in either case.

Orbit.	Surf.	Sol	LMST	Assign.	FWD	RET
ERO	M20	125	8	RET	5	400
MVN	SFR	125	11	TOTAL	2	200
DFE	M20	125	22	FWD	1	0

network of Mars orbital craft, these communication opportunities dictate the available *relay passes*. For each such window, MARMLADE estimates the bandwidth available for data sent from Earth to Mars (uplink) or from Mars to Earth (downlink), using the geometry of the Earth–Mars connection. Allocations are computed based on parameters and assumptions (e.g., the data requirements of a given surface team) supplied by mission planners.

Data

Each optimized pass allocation consists of a table of pass assignments [see the data schema (Table 1)], which includes several nominal columns: *orbital asset* describes which hardware would broker the relay pass: either one of four Mars orbiters, or direct-from-Earth (DFE) X-band communication. Similarly, *surface asset* describes which robot on Mars would receive this much dedicated bandwidth. Finally, *allocation assignment* describes the type of pass, as one of RET (return from Mars), TOTAL (non-time-critical return), or FWD (forward to Mars).

Date-time columns provide pass timing and duration in Marstime (LMST) and Earthtime (UTC). The continuous *forward volume (Mb)* and *return volume (Mb)* variables describe the bandwidth, in megabits, which could be carried by the pass in uplink or in downlink mode, respectively. Several internal optimization choices are also exposed, including *Split Pass?*, a Boolean indicating whether this relay pass has actually been divided between two surface assets; each one gets its own record.

On each Mars day (i.e., a numbered *sol*), a closed operational cycle consists of an uplink pass (where instructions go to the surface asset on Mars), then an execution window, followed by a downlink pass (where the asset returns its data for the day’s operations). Our stakeholders seek to maximize *operational efficiency*

TABLE 2. Listing of tasks reported by the relay and surface team leads, motivating the MarsIPAN schedule design.

	Diagnosis
T1a	Locating passes.
T1b	Seeing a missed uplink.
T1c	Changing scenario assumptions.
	Negotiation
T2a	Highlighting a bandwidth need.
T2b	Characterizing a staffing window.
T2c	Sharing data with other scheduling tools.
	Reflection
T3a	Modifying the tradespace of assumptions.
T3b	Investing trust in the allocation algorithm.

by completing as many operational cycles as possible for each surface asset.

Users

At the outset of our design study, with a 10-week commitment from the allocation team, we set out to *identify* other stakeholders and to *characterize* the visualization problem. While Sedlmair et al.⁹ rightly characterize the identification of stakeholders as a precondition to problem characterization, we were motivated to conduct both in parallel due to the complexities of connection, translation, and project review at JPL. We believed that shared bandwidth allocations were not being adopted due to usability problems, given that the existing visual representation of an allocation schedule was hand-authored in Excel. As nonaerospace experts, once we had created a system capable of locating passes (T1a, Table 2), the allocation team taught us how to see a missed uplink (T1b), and that each allocation depends on different assumptions, which vary across threat scenarios (T1c).

As we spoke with relay and surface experts to discover their needs and to demonstrate the MarsIPAN prototype's capabilities, we learned about their professional duties of translation and analysis. Since we were still engaged in task wrangling,¹⁰ we realized that diagnosis alone was insufficient, and both negotiation and reflection tasks were vital to the success of MSR. MarsIPAN would need to support experts' evidence-based arguments in presentations to mission leadership and other decision-makers. In order to capture this extended cast of stakeholders, which included senior JPL experts

who engaged with the MarsIPAN system but did not have time to be extensively interviewed by our team, we developed the knowledge transit diagram (see Figure 1).

Each team on MSR operates their own infrastructure, including vehicles and other information systems. The *orbital relay team* operates communications satellites around Mars and is responsible for allocating pass bandwidth to each of the surface missions. Each *surface robotics team* consumes that bandwidth through uplink (data to Mars) and downlink (data from Mars) passes in order to perform engineering and science operations each Mars day (see Figure 2). During the planning phase of MSR, both relay and surface teams regularly analyze their own capabilities to establish whether they have enough support to carry out planned operations, within a safe margin of resources. Team leads are domain experts who are skilled at translation, making their capabilities and requirements clear to other teams that share the same resources (i.e., time and bandwidth).

Mission analysts determine how to fit surface activities, such as operating the vehicle and downlinking its science data, into the available relay pass bandwidth. Surface teams model their bandwidth need for each activity using historical data. Each surface team wants to collect both decisional data (which is required for the vehicle to operate) and scientific data on each Mars day.

Each surface team wants to close its *operational cycle* on each Mars day. That is, each spacecraft must receive uplink data before the execution window on Mars (i.e., when the surface asset has enough sunlight to drive around and use its instruments), based on data returned to Earth after the previous execution window. Downlinks include sensor data from the surface vehicle, which is incorporated into planning for the next sol. Since planning itself is time-consuming, it is possible to miss out on an execution window because the uplink opportunity occurred too soon after downlink. When this happens, a sol is referred to as *restricted*, and the spacecraft cannot be commanded until the following day, missing critical opportunities to act within the finite mission lifetime.

Orbital Relay Team

The relay operations team provides relay capabilities across the mission duration, including both the orbiter's power requirements and physical line-of-sight constraints between the operations site on Mars and the communications antenna. Each "relay pass" is actually a promise made by relay about their ability to

provide bandwidth over that span of time. The relay lead wants to maintain the health of relay passes in adversarial scenarios, ensuring that the available communications hardware satisfies mission requirements even if something goes wrong. They work closely with the allocation team to make sure that relay passes are calculated based on recent data.

Surface Robotics Team

Each surface asset has its own operations lead. The surface operations team wants to know when their team will have data and be able to plan. They are concerned with coordinating different phases of the mission, which includes reasoning about the schedule in terms of Earthtime working hours. In the past, staffing schedules were created according to 25-h Mars days, leading to excessive overtime for long periods that would demoralize staff and are not sustainable for the mission. Each surface team lead reports to the MSR operations lead, who needs to make sure that relay opportunities are distributed fairly between teams.

Allocation Design Team

Our visualization team was itself embedded with a third team at NASA/JPL, which was responsible for implementing the allocation algorithm in order to reduce conflicts over pass trading during MSR. As such, the allocation team recruited us to improve the usability of allocations, which were produced in a bare CSV format that was difficult for other teams to read. Because we worked closely with the allocation lead to understand their data, meeting on a weekly basis to evaluate MarsIPAN prototypes for several months, we elected to evaluate the system's usability with other NASA/JPL experts who were interested enough in the allocation team's claims to look at their data.

Tasks

The MarsIPAN schedule enables experts to visually identify "problem spots" in pass allocations, characterized by missed operational cycles. Experts use the MarsIPAN schedule to evaluate optimal pass allocations and identify the sources of various problems. We break down this process into diagnosis tasks, negotiation tasks, and reflective tasks.

Diagnosis tasks (T1a–c) include getting oriented in the pass allocation, localize problems to certain weeks or certain days of operation, and contextualizing within the constraints of the scenario. Stakeholders wanted to look closely at allocations to establish whether they really are optimal, or if opportunities are

being missed (i.e., require additional contextual knowledge to identify).

Negotiation tasks (T2a–c) include comparing allocations side-by-side (by opening two windows), taking screenshots of allocations in order to share notes, and excerpting the pass allocation for further analysis in team-specific tools. A large tradespace of optimal allocations is possible, depending on unforeseeable operational constraints (e.g., whether a legacy orbiter survives) and decisions made by other teams.

Reflection tasks (T3a–b) consist of evaluating the allocation optimization strategy as a whole, as to whether it actually reduces the burden of manually allocating passes across teams. The MARMLADE designers also sought to engage fellow experts in conversations about allocation quality and the potential impacts of various threat scenarios.

VISUALIZATION SYSTEM

The MarsIPAN interface is made up of three linked widgets. First, the Navigation track on the left (see Figure 3 ① and Figure 4 ④) overviews operational cycles during the mission. Second, the main Schedule panel [see Figure 3 ② and Figure 4 ③] displays the relay passes over three Mars days. It can group passes either by their assigned surface vehicle or their source orbiter, and displays both Marstime and Earthtime, either of which can be the primary x-axis. Third, the Pass Wrangler on top of the schedule [see Figure 3 ④ and Figure 4 ①] allows analysts to select and filter passes based on data volume, assignment, and other characteristics.

These three MarsIPAN widgets also correspond to three layers in a hierarchy of abstraction (see Figure 5), following a typical progressive disclosure pattern and allowing users to interact with the data at the appropriate level of granularity to match their tasks. The three widgets are described below, followed by a description of how selections are shared between these widgets. We defer discussion of the highest level of abstraction, the *tradespace* of all possible allocations, to the "Evaluation" section.

Navigation Track

The Navigation track occupies the second-highest level of abstraction (see Figure 5), providing a graphical *barcode*, which summarizes the operational cycles within a given allocation. It is split into two zoom levels: the dense COMP track, and the colorful M20/SRL/SFR tracks, respectively. Split minimaps are also used in genome browsers,²⁰ where analysts are interested

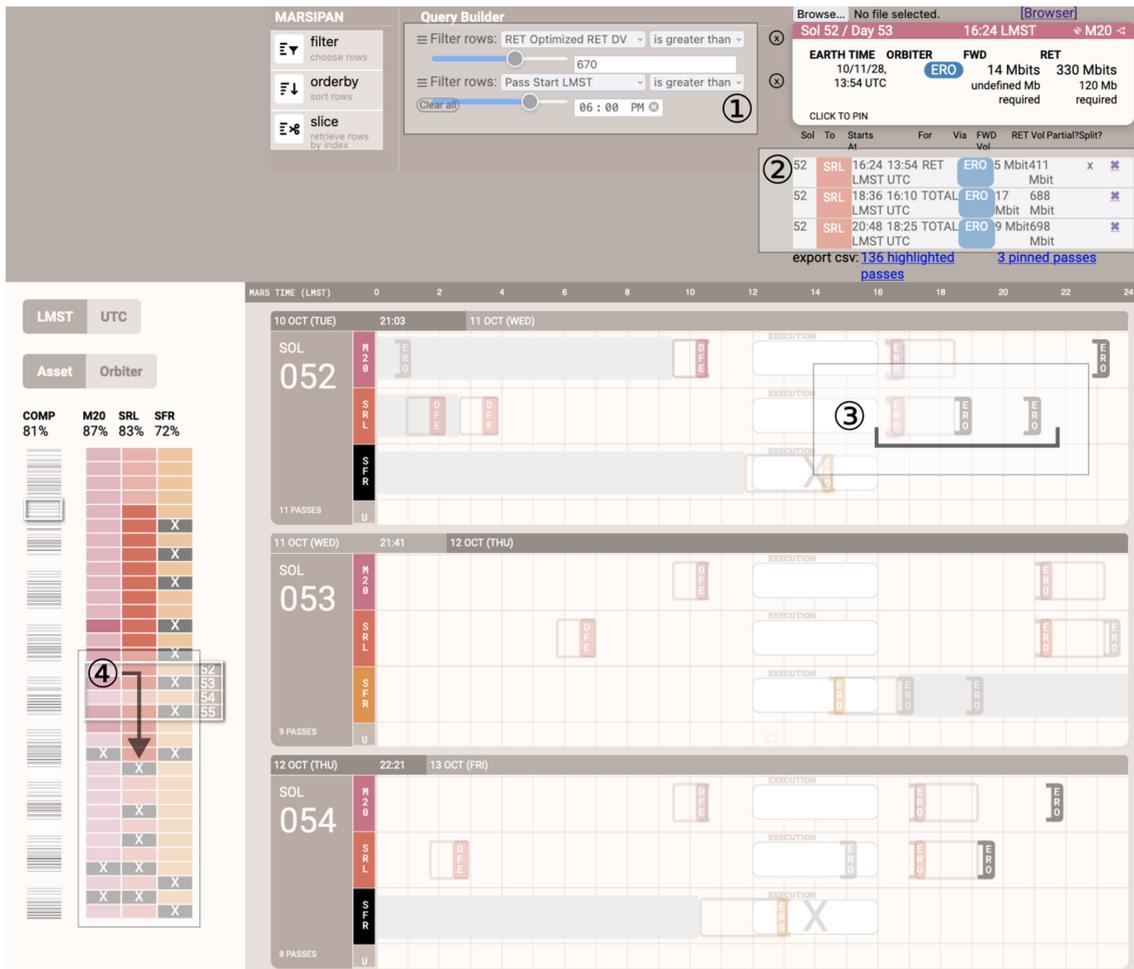


FIGURE 4. Use case demonstrating the query builder. In this figure, we consider a subset of passes in the evening on Mars, after the execution window.

Mars 2020 rover (M20), the Sample Retrieval Lander (SRL), or the Surface Fetch Rover (SFR). Split passes can be assigned to two teams and appear in parallel on their tracks. The “pass pattern” describes the timing of all passes supplied by the same orbiter; for example, the ERO follows a regular six-pass cadence each Mars day.

Pass Wrangler

At the lowest level of abstraction (see Figure 5), the Pass Wrangler sits across the top of the interface. It is used to filter data within an allocation, allowing users to locate and save passes with interesting characteristics. It consists of three parts: the query builder (left), the pass detail panel (top-right), and the bookmarked passes (bottom-right).

Queries are expressed in terms of filtering, sorting, and slicing operations, which may be exported

to a CSV file. Filtering operations use sliders, UTC and LMST timestamps, and dropdown selections appropriate to each column’s data type. A small button allows individual filters to be removed. The surface team lead was able to express a Crossfilter-like query through this interface, producing a selection of passes which was highlighted throughout the schedule.

The pass detail panel is activated when the expert hovers their mouse over any pass in the schedule, displaying properties of that pass from the *ground* truth allocation, following Table 1. These include exact start times in UTC and LMST, and its potential uplink or downlink bandwidth, compared to the assigned surface asset’s data volume requirements.

Clicking on a pass adds it to the bookmark panel beneath the pass info box, so that its details can be

Abstraction

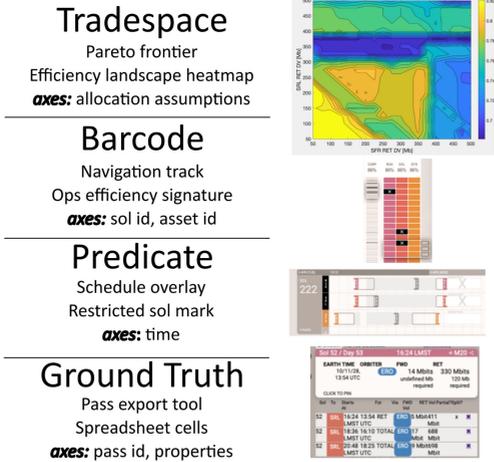


FIGURE 5. Thumbnails of three MarsIPAN widgets—the Pass Wrangler, the Schedule panel, and the Navigation track—representing successive levels of abstraction over the relay pass allocation. Finally, the tradespace figure (reproduced from Donitz et al.¹⁹ with permission) represents one slice of all possible allocations.

compared with other passes. A small button allows passes to be unbookmarked. Nonempty query selections applied to the set of passes are represented by higher opacity. Both selected passes and bookmarked passes can be exported to CSV for subsequent analysis.

Selections

All three widgets in the MarsIPAN schedule are bidirectionally linked according to a hierarchical relationship of ascending abstraction. The Navigation track is used to scroll up and down the Schedule view, on a day-by-day basis. Likewise, the Schedule view can be clicked on to save passes into the Pass Wrangler, which can then be used to keep track of places in the Schedule panel. Query selections made in the Pass Wrangler propagate to both the Schedule panel and the Navigation track, allowing large-scale patterns to be observed and inspected simultaneously.

A modal button (LMST/UTC, above the navigation track) switches the schedule and navigation track between Earth days and Mars days (which span almost 25 Earth hours). A second modal button (Asset/Orbiter, above the navigation track) flips the schedule panel’s facets (which orbiter provided the pass) and the labels (which surface team receives the pass), and changes the navigation track into a pass

pattern overview; each orbiter displays all of its passes on its own layer.

IMPLEMENTATION

MarsIPAN was built in a reactive visualization notebook, which gave us flexibility to respond to changing user requirements, and was straightforward to deploy on a GitHub Enterprise instance. As the application code belongs to NASA/JPL, we are unable to share the repository itself.

We used the Observable runtime^a as a reactive prototyping framework for the client-side Javascript application. Observable integrates both HTML-based and SVG-based visualization widgets, typically written using D3.js, into larger applications with persistent selections. The MarsIPAN query builder is derived from a demo of the Arquero query library, adding appropriate form inputs for both nominal and date-time fields.^b

USAGE SCENARIOS

MarsIPAN is used by different teams across MSR to perform the analysis tasks appropriate to their specific rolea. We provide two click-through scenarios in this section to illustrate how experts use the schedule analysis interface to explain, understand, and reason about one scheduling algorithm contingency.

Scenario: Single Operational Cycle

The *surface lead* for SRL wants to know how much data they will receive from operations on Sol 52, and when it will arrive (task T2a) (see Figure 4).

① The surface lead creates a query for passes with high return volumes (“RET Optimized RET DV”) in the Mars evening (“Pass Start LMST”), as soon as the vehicle has finished executing its instructions.

② They identify three return passes (closing brackets), all supplied by ERO. The first of these is the decisional return pass (marked in red and with a latency tail), which is split between SRL and M20. They inspect each pass by hovering over it, once in the SRL track (which has a return volume of 411 Mbits) and once in the M20 track (which carries 330 Mbits), verifying that both missions get enough return volume to begin operations planning (i.e., their decisional return needs are met), and that the total pass bandwidth is as expected.

^aOpen source. [Online]. Available: github.com/observablehq/runtime

^bOpen source. [Online]. Available: observablehq.com/@observablehq/data-wrangler

③ The surface lead performs a time conversion from UTC to PST to determine that after latency, decisional return arrives around 7 A.M. PST, which will make staffing for Sol 52 easy.

④ To see if the pattern of high-quality passes will continue, they look at the left-hand side navigation track. Sliding down into the future, they observe that in 5–10 sols, the mission will run into a series of days where passes do not complete. They interpret this to be due to the sliding Earth–Mars time window, which eventually encroaches on the Earthtime work day window.

Scenario: Split Pass

To diagnose a missed uplink (Task T1b), the *relay lead* discusses the availability of split passes (see Figure 3) on legacy orbiters with an algorithm designer on the *allocation design team* in order to *maintain relay pass health* and identify relevant tradespace assumptions (Task T1c). They are interested in the new “FDMA split pass” capability, which allows two surface teams to efficiently share a pass by “splitting” its bandwidth between distinct sets of frequencies.

① The relay lead looks at the operational efficiency figures, seeing that M20’s operational efficiency is very high (98%). They identify Sol 80 as one of the missed operational cycles, by observing the strike mark in the bar code, and scroll to that point in the schedule.

② The algorithm designer explains the algorithm’s decision: required data did not arrive until 5:00 LMST, which is too close to the execution window (white rectangle beginning 12:00 LMST) for planning (which takes 8 h, shaded rectangle ending 13:00 LMST) to be completed.

③ The relay lead sees on Sol 79 that a single Mars Reconnaissance Orbiter (MRO) pass is available in the evening and points out that M20 could have down-linked much sooner if it was split. Since there is enough data volume to split the pass, they conclude that pass splitting was not turned on for MRO in this allocation, just for ERO. This knowledge about the algorithm is not actually present in metadata and must have come from direct conversations with the allocation team.

④ The algorithm designer suggests that M20 would not be able to complete its required return anyway, since only 143 Mbits would be left over after SRL consumes 175 of the 318 Mbits available. The relay lead responds that SRL’s needs will likely be reduced from 175 to 150 Mbits. Their implication is that the M20 operational cycle on Sol 80 would be closed if this parameter to the scheduling algorithm were updated.

EVALUATION

The MarsIPAN prototype has been released directly to stakeholders and is being used in production by the various stakeholder communities across the MSR mission to inform their mission planning decisions. To assess the impact of the deployed application, we formally evaluated the MarsIPAN schedule prototype using a think-aloud protocol,²¹ asking participants to locate and diagnose an operational efficiency loss by interacting with real data using the MarsIPAN system, in tasks co-designed with stakeholders. Workflow was not constrained during testing, and stakeholders were allowed to freely explore the system. After participants completed this task, reflections were collected using retrospective semistructured interviews.²² The protocol was performed with participants from two user groups—the relay operations team lead, and one of the surface operations team leads,—for one hour each over video-conferencing, while the expert screen-shared the application. For analysis, data were manually transcribed and then coded for emergent theme analysis,²³ resulting in the tasks described in Table 2.

We asked experts to diagnose missed operational cycles in a single allocation, representing one of several actual MSR communications scenarios under analysis at the time. We also asked the expert to walk us through what kinds of questions they would ask about these data, whether that was through the interface or of another team.

Experts demonstrated deep familiarity with the data, quickly becoming oriented in the context of a scenario and directing us toward high-level analysis problems. For example, both experts asked how the structure of the allocation would be impacted by changes to the upstream scenario. Experts navigated the allocation and pointed out problems to us using each of the MarsIPAN views, moving fluidly between low-level and high-level perspectives,⁷ which enabled new tasks to emerge; in addition to the diagnosis we had asked for, the experts also freely explained pass negotiations and their own process of reflection.

Diagnosis

T1a—Locating Passes

The relay expert located a pass labeled ERO in the schedule panel. They followed successive passes across the timeline to situate themselves on the sixth pass provided by that overflight of ERO, becoming oriented “in the pass pattern.” From this position, they mentioned that multiple MRO passes should follow it: “Yeah, I only see one MRO, so where are the others?”

There should be at least two opportunities, if not three or four.” We followed up with the allocation design lead to confirm this, and they concluded that the second MRO pass had been dropped by the allocation algorithm due to its low data volume.

T1b—Seeing a missed uplink

The relay lead characterized flaws in the MarsIPAN “ground-truth” data arising from coarse-grained upstream assumptions. They navigated to Sol 193 using the track, seeing that it was struck, and located the critical return pass at 14:44 LMST in the schedule. They moused over the pass to see that its return volume was 33 Mb, much lower than the 55 Mb they expected to see for this type of pass, characterizing the difference as “conservatism upstream” because a low-elevation pass could be blocked by mountain ranges between the orbiter and the operations site; however, in practice, the worst-case assumption was not often borne out. Therefore, they reasoned that the pass should be split between its assigned surface vehicle and the struck one so that both could operate the next Mars day.

T1c—Changing an Assumption

The surface expert asked a question about the impact of a 24-h staffing scenario while they were looking at a specific Mars day in the version of the allocation with 18-h staffing. We had both allocations available and prompted them to load the other allocation, but this caused them to lose their place in the schedule from which to make the comparison. While they suggested implementing a head-to-head comparison, we realized later that multiple instances of MarsIPAN can be loaded in separate browser windows, and believe that experts adopted this practice quickly.

Negotiation

T2a—Diagnosing a Missed Uplink

The relay expert used MarsIPAN to investigate a pattern of missed uplinks in a pass allocation, observing that the forward pass overlapped the execution window for several Mars days at a time: “For whatever reason they can’t command until after their execution window [on Sol 5]. I think this issue is more on Sol 6 than Sol 7. And the ERO pass that would be available is right in the middle of the execution window [on Sol 9].”

The specific characteristics of MRO, which provided the evening pass on Sol 9 at 19:15 LMST, meant this pass could not be split between both surface missions (so that one became restricted). “M20 and SRL

should probably be alternating back and forth [...]” in terms of which one was restricted every other Mars day. While the low quality of that evening pass was inherent to the pass pattern—the expert verified that it repeated every 40 Mars days, by scrolling through the navigation track,—they suggested that the burden had not been shared fairly between surface teams.

T2b—Communicating Team Needs

The surface lead noticed that the DFE FWD on Sol 7 had been allocated during the execution window. They switched the schedule view from Marstime to Earthtime, identifying the decisional RET arriving at 5:30 A.M. PST, and concluding that the nominal staffing start at 6 A.M. had led to the planning overrun. They were satisfied that actual staffing would not miss this sol, because during the first two weeks, their team intended to staff on Marstime already.

Just as the relay lead had identified that the second MRO pass on Sol 7 was missing, the surface lead noticed this issue next. We communicated that the algorithm had dropped this pass for its low volume, so the surface lead decided to ask the algorithm designer about that choice. “Great for asking questions about their algorithm. [...] I’d use the information that’s here and call it out. That orients it toward the problem or concern we’re talking through.”

T2c—Complementing Team-Specific Tools

The surface lead characterized the MarsIPAN schedule as being suitable for different kinds of tasks than their team’s existing planning tool. In terms of layout, the surface planning tool does not split days vertically, so operational cycles are never broken vertically across the schedule; instead, it is possible to scroll horizontally through the days in order to establish the fine-grained relationship between Marstime and Earthtime. Although their system can interface with state timelines, which are used to orchestrate complex behaviors, it lacks the navigation and overview capabilities of MarsIPAN:

“The planning tool consumes [temporal event] data all the time. We have the ability to add state timelines underneath, and write constraints to check for conflicts. MarsIPAN is the tool we [would] use to sanity check that what we’re saying is correct. More of a strategic planning tool and a development planning tool.” Strategic planning is when teams find days they won’t be able to operate, in order to determine their baseline operational efficiency in the reference scenario (where nothing has gone wrong). Development planning lets teams understand which capabilities in the available orbital assets and surface assets are

most desirable and should be prioritized for development by the time the mission launches.

Reflection

T3a—Getting Oriented in a Tradespace

The surface expert used the navigation panel to find a cluster of strike marks, characterizing that region of the schedule as holding special concern: “It looks like I can see where the clustering [of missed decisional cycles] is that I should be concerned. Like this chunk is in the prime mission.” They described their next steps in this scenario as “trading slack” with another team, either surface or relay, who would give up some of their marginal bandwidth in recognition of the other team’s critical need. While they wanted to use MarsIPAN visuals to help make this argument, they also noticed that other forms of “slack” were embedded in assumptions of the allocation algorithm and wanted to explore these tradeoffs to support their case to the other team.

T3b—Evaluating Algorithm Assumptions and Decisions

The relay expert asked for an option to recompute the allocation from within MarsIPAN with updated assumptions; or failing that, to manually override the allocation of a single pass, correcting a fault introduced by the algorithm. (The surface lead had a contrasting approach to allocations, regarding them as ephemeral predictions to be used in active negotiations about the tradespace parameters.)

“I don’t fully trust the algorithm to do a better job than me, and I can find opportunities. [...] For these sections I would let it do its thing and then come [scrolls] down here and fix it.”

The relay expert focused their assessment of the algorithmic allocation on difficult sections of the schedule, emerging with a stronger understanding of which assumptions entered the allocation, and where these might be flawed. “It makes [assumptions] much easier to work with and digest, rather than something that’s hidden away.”

They expressed an intent to continue engaging the systems engineering team directly, armed with a deeper understanding of their work, in order to improve the quality of the scheduling algorithm.

DISCUSSION

The design study method positions the visualization artifact in service to stakeholder needs, which not only address an appropriate formulation of the problem but also support knowledge transfer across

multiple disciplines in an organization. The MarsIPAN design study produced an artifact that serves stakeholder needs organized around four cohesive levels of abstraction (see Figure 5) in relay pass allocation.

At the *ground* level, the “relay pass” enables relay experts to convey relay capabilities while “folding away” technical details of relay operation. At the next highest level of abstraction, the decisional cycle *predicate* characterizes the relationship between uplink and downlink decisions with respect to the execution window on Mars, and the staffing window on Earth. Aggregating the decisional cycles into a minimap produces the *barcode*, a summary representation of operations health over time. At the highest level of abstraction, the *tradespace* of optimal pass allocations is characterized by differences in operational efficiency across various possible scenarios.

The MarsIPAN schedule enables us to study how project teams conduct analyses at different levels of abstraction, reflecting shared concerns in the context of different areas of responsibility and expertise. While we produced direct evidence of analysis, and only indirect evidence of coordination, we employ seamful design²⁴ to articulate how it is possible for teams that deal with wildly different forms of complexity to successfully negotiate with each other. We identified two concerns shared by all experts: (1) the need for explainability, with respect to each discipline, and (2) the complexity of translating satisfying explanations from one discipline to another.

Explaining Bandwidth Allocation Decisions

Before MarsIPAN, analysts could see high-level abstractions, such as operational efficiency, but they could not understand the reasons driving those statistics, nor explain how changes to the underlying pass allocation algorithm caused (or did not cause) changes to the “key metric” of operational efficiency. MarsIPAN provides a cohesive set of linked views which experts used to unpack, explain, and ultimately justify the operational efficiency metric.²⁵

MarsIPAN increased trust in the diagnosis process around allocations, and through so doing, changed the negotiation processes around bandwidth, helping experts reflect on MSR scenarios. Because they could see allocations clearly, users were able to determine when single passes were allocated badly, and motivated changes to the pass allocation algorithm. Similarly, since users were able to uncover examples where split passes were not fairly allocated across multiple missions, MarsIPAN introduced a shared

common ground toward redesigning the way that split passes were organized across missions. Because users could find places where operational cycles were not closed because assumptions were too tight, users were able to find instances where relaxing assumptions around the scheduling window led to substantially different allocations, driving discussion around what those correct assumptions need to be within and across teams.

Some missed operational cycles are essential, and others represent lost opportunities Experts distinguished between circumstances where the pass allocation correctly indicated a missed operational cycle, and others where it had failed. Experts from different teams made different suggestions as to how and why the model had failed, based on their own deep knowledge of the scheduling constraints. These relationships form a basis for algorithmic trust, using the visualization.²⁶ MarsIPAN enables experts with a variety of domain backgrounds to interpret and control the results of the MARMLADE algorithm.

Algorithmic Explainability

Relay passes are abstractions which our experts readily contested. Experts sought to complicate assumptions when and where those changes might impact decision-making and ultimately leveraged their own knowledge of the MSR project to ask other stakeholders for changes to the mission design parameters. Because of well-chosen visual representations, MarsIPAN became a “transparent” interface to the pass allocation data.

Schedule optimization ranks candidate solutions by their operational efficiency, a metric defined in terms of (missed) decisional cycles. We discovered that JPL experts treat the decisional cycle as a seam²⁴ into which the relay passes could disappear safely. Relay passes are yet another seam, concealing the complexity of relay operations from nonrelay experts. The MarsIPAN schedule makes operational cycles available to annotation and discussion for the first time, in terms of individual relay pass opportunities. While surface teams need to trust that they have gotten as much bandwidth as possible, given everyone’s needs, mission planners need to agree that the assignments are well chosen and contain a reasonable margin of safety.

MarsIPAN supports two forms of explainability: (1) as part of existing diagnosis workflows, and (2) in support of negotiation with other teams. Incorporating the MarsIPAN schedule into allocation analysis increased the algorithm’s *explainability*, which is

desirable in resource allocation problems. Our design study intervention empowered experts to articulate why and how they would incorporate schedule optimization into existing negotiations between teams, not just (1) to avoid making trivial decisions over and over as other parameters of the mission were changed, but also (2) to clearly illustrate reports about mission capabilities and resilience.

Knowledge Transit Diagrams

Experts addressed scheduling problems by creating common ground that encapsulates their expertise, in order to collectively develop a shared understanding of their respective goals (which were sometimes at odds). Shared visualization systems enable these experts to exploit shared abstractions—that is, knowledge work that enables their distinct ways of knowing to form a singular coherent narrative. In order to demonstrate the mediated exchange of expertise at NASA/JPL, we introduce the *knowledge transit diagram* (see Figure 1), an artifact that theoretically encodes multidisciplinary across diverse stakeholder groups.

Stations of the diagram (white circles) correspond to stakeholder roles, containing a tight “reflective loop” enabled by a precise vocabulary of concepts and data frames. Once knowledge is produced inside of a station, it must circulate to other stations through conversations and reports. The spaces between stations correspond to observable *seams* wherever they are bridged by datasets, visualization artifacts, and other representations of abstraction.²⁴

Lines (colorful ribbons) of the diagram are domains of knowledge containing extensively interrelated sets of abstractions and metaphors, whose combined effect is to facilitate clear and rapid discussions concerning systems that fall within the domain. The absence of these compatible layers of abstraction is typically experienced as a *data friction*,²⁷ whereas their presence is characteristic of domain expertise²⁸ and disciplinary formations.²⁹

The knowledge transit diagram leads us to characterize design studies as a form of workplace ethnography, and a site where teams may collaboratively build a shared language and set of abstractions³⁰ to mount a collective campaign that bridges disciplines.¹⁷ Our team worked with Orbital Relay (teal) to illustrate optimal pass allocations in a reproducible way. By creating technology to mediate the negotiation around what counts as optimal, we also bridged another seam and empowered the adjacent stakeholder in surface robotics team to negotiate

their staffing needs. They chose to join our design study because we were open to addressing human factors in operations scheduling at the tradespace level and able to ground that conversation in the underlying pass data (see Figure 5).

Our study informs broader practices of visual encoding and interaction around generalizable tasks and abstractions. Through the knowledge transit diagram, we open the black box of interactions between MSR teams within NASA/JPL to reveal complex forms of sociotechnical delegation that perform continuous repair at various abstract boundaries within this interdisciplinary organization.²⁹ Design study practitioners should consider returning to the casting stage well after preliminary work,⁹ in order to map the structure of interfaces between individuals, teams, and disciplinary commitments that characterizes design work itself.⁸

CONCLUSION

Visualization practitioners are uniquely positioned to gather and circulate knowledge among teams that are committed to engaging with each other, yet do not share a common language. This article reflects on a two-year design study supported by multiple stakeholder groups involved in multilateral robotics operations and scheduling at NASA/JPL. MarsIPAN facilitated productive conversations between collaborators, both within and outside of the design study context, by transforming a set of disparate technical and operational threats into shared knowledge about project capabilities and requirements. Visualization researchers have much to gain by supporting practitioners who create knowledge infrastructure, whether they are part of our professional field, or experts with whom we collaborate.

ACKNOWLEDGMENTS

This work was supported by the Jet Propulsion Laboratory, California Institute of Technology, through a contract with the National Aeronautics and Space Administration under Grant 80NM0018D0004.

The authors would like to thank their anonymous JPL experts for their insight and contributions to this work; Danyel Fisher and Jack V. Maydan for additional conversations; and the co-organizers of the Data to Discovery program, including Santiago V. Lombeyda, Hillary Mushkin, and Maggie Hendrie.

This work involved human subjects or animals in its research. Approval of all ethical and experimental

procedures and protocols was granted by the California Institute of Technology under Application No. IR24-0461, and performed in line with the Improving User Experiences in Making Decisions About Science protocol.

REFERENCES

1. J. Jo, J. Huh, J. Park, B. Kim, and J. Seo, "LiveGantt: Interactively visualizing a large manufacturing schedule," *IEEE Trans. Vis. Comput. Graphics*, vol. 20, no. 12, pp. 2329–2338, Dec. 2014.
2. T. L. Lei and R. L. Church, "Mapping transit-based access: Integrating GIS, routes and schedules," *Int. J. Geograph. Inf. Sci.*, vol. 24, no. 2, pp. 283–304, Feb. 2010, doi: [10.1080/13658810902835404](https://doi.org/10.1080/13658810902835404).
3. B. K. Muirhead, A. Nicholas, and J. Umland, "Mars Sample Return Mission Concept Status," in *Proc. IEEE Aerosp. Conf.*, Mar. 2020, pp. 1–8.
4. J. Schier, J. Rush, P. Vrotsos, and W. Williams, "Space communication architecture supporting exploration and science: Plans & studies for 2010–2030," in *Proc. 1st Space Exploration Conf.: Continuing Voyage Discov.*, 2005, Art. no. 2517.
5. H. A. Simon, "Rational choice and the structure of the environment," *Psychol. Rev.*, vol. 63, no. 2, 1956, Art. no. 129.
6. U. Hinrichs, S. Forlini, and B. Moynihan, "In defense of sandcastles: Research thinking through visualization in digital humanities," *Digit. Scholarship Humanities*, vol. 34, no. Supplement_1, pp. i80–i99, Dec. 2019.
7. A. Diehl et al., "An analysis of the interplay and mutual benefits of grounded theory and visualization," *IEEE Trans. Vis. Comput. Graphics*, early access, Sep. 02, 2024, doi: [10.1109/TVCG.2024.3452985](https://doi.org/10.1109/TVCG.2024.3452985).
8. J. T. Otto and S. Davidoff, "Visualization artifacts are boundary objects," in *Proc. 10th IEEE Eval. Beyond - Methodological Approaches Visualization*, Aug. 2024, pp. 81–88.
9. M. Sedlmair, M. Meyer, and T. Munzner, "Design study methodology: Reflections from the trenches and the stacks," *IEEE Trans. Vis. Comput. Graphics*, vol. 18, no. 12, pp. 2431–2440, Dec. 2012.
10. A. Crisan and T. Munzner, "Uncovering data landscapes through data reconnaissance and task wrangling," in *Proc. IEEE Vis. Conf.*, Vancouver, BC, Canada, Oct. 2019, pp. 46–50.
11. J. J. Marquez, G. Pyrzak, S. Hashemi, K. McMillin, and J. Medwid, "Supporting real-time operations and execution through timeline and scheduling aids," in *Proc. 43rd Int. Conf. Environ. Syst.*, Jul. 2013, Art. no. 3519.

12. A. Rind, T. Lammarsch, W. Aigner, B. Alsallakh, and S. Miksch, "TimeBench: A data model and software library for visual analytics of time-oriented data," *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 12, pp. 2247–2256, Dec. 2013.
13. S. S. Bae, F. Rossi, J. V. Hook, S. Davidoff, and K.-L. Ma, "A visual analytics approach to debugging cooperative, autonomous multi-robot systems' worldviews," in *Proc. IEEE Conf. Vis. Anal. Sci. Technol.*, Oct. 2020, pp. 24–35.
14. J. Kim et al., "Cobi: A community-informed conference scheduling tool," in *Proc. 26th Annu. ACM Symp. User Interface Softw. Technol.*, New York, NY, USA, Oct. 2013, pp. 173–182.
15. J. Bellien, B. Cardoen, and E. Demeulemeester, "Improving workforce scheduling of aircraft line maintenance at Sabena technics," *Interfaces*, vol. 42, no. 4, pp. 352–364, Aug. 2012.
16. M. Saint-Guillain, J. Gibaszek, T. Vaquero, and S. Chien, "Romie: A domain-independent tool for computer-aided robust operations management," *Eng. Appl. Artif. Intell.*, vol. 111, May 2022, Art. no. 104801.
17. V. Riboulet, P. Marin, and J.-C. Leon, "Towards a new set of tools for a collaborative design environment," in *Proc. 7th Int. Conf. Comput. Supported Cooperative Work Des.*, Sep. 2002, pp. 128–133.
18. M. Brehmer, S. Carpendale, B. Lee, and M. Tory, "Pre-design empiricism for information visualization: Scenarios, methods, and challenges," in *Proc. 5th Workshop Beyond Time Errors: Novel Eval. Methods Visual.*, New York, NY, USA, Nov. 2014, pp. 147–151.
19. B. P. S. Donitz, A. M. Didion, A. K. Nicholas, T. J. Voss, and C. H. Lee, "Assessing relay communications for Mars Sample Return Surface Mission Concepts," in *Proc. IEEE Aerosp. Conf.*, Mar. 2021, pp. 1–10.
20. P. C. Sabeti et al., "Positive natural selection in the human lineage," *Science*, vol. 312, no. 5780, pp. 1614–1620, Jun. 2006.
21. O. Alhadreti and P. Mayhew, "Rethinking thinking aloud: A comparison of three think-aloud protocols," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Montreal, QC, Canada, Apr. 2018, pp. 1–12.
22. A. Blandford, "Semi-structured qualitative studies," in *Proc. Encyclopedia Hum.-Comput. Interact.*, 2nd ed. Aarhus, Denmark: Interact. Des. Found., 2013.
23. K. Charmaz, "Grounded theory as an emergent method," in *Handbook of Emergent Methods*. New York, NY, USA: Guilford Press, 2008, pp. 155–170.
24. S. Inman and D. Ribes, "'Beautiful seams': Strategic revelations and concealments," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, New York, NY, USA, May 2019, pp. 1–14.
25. M. Conlen et al., "Towards design principles for visual analytics in operations contexts," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Montreal, QC, Canada, Apr. 2018, pp. 1–7.
26. E. Beauxis-Aussalet et al., "The role of interactive visualization in fostering trust in AI," *IEEE Comput. Graphics Appl.*, vol. 41, no. 6, pp. 7–12, Nov./Dec. 2021.
27. G. Panagiotidou, J. Poblome, J. Aerts, and A. Vande Moere, "Designing a data visualisation for interdisciplinary scientists. how to transparently convey data frictions?," *Comput. Supported Cooperative Work*, vol. 31, no. 4, pp. 633–667, Dec. 2022.
28. J. Li, E. Rawn, J. Ritchie, J. Tran O'Leary, and S. Follmer, "Beyond the artifact: Power as a lens for creativity support tools," in *Proc. 36th Annu. ACM Symp. User Interface Softw. Technol.*, San Francisco, CA, USA, Oct. 2023, pp. 1–15.
29. D. Ribes, A. S. Hoffman, S. C. Slota, and G. C. Bowker, "The logic of domains," *Social Stud. Sci.*, vol. 49, no. 3, pp. 281–309, Jun. 2019.
30. J. Rezwana and M. L. Maher, "Designing creative AI partners with COFI: A framework for modeling interaction in Human-AI co-creative systems," *ACM Trans. Comput.-Hum. Interact.*, vol. 30, no. 5, pp. 67:1–67:28, Sep. 2023.

JASMINE T. OTTO is currently with the University of California, Santa Cruz, CA, 95064, USA. She works closely with domain experts to develop novel intelligent interfaces supporting professional vision and collaborative sensemaking practices. She is the corresponding author of this article. Contact her at jtotto@ucsc.edu.

MALIKA KHURANA is currently with *The New York Times*, New York, NY, 10018, USA. She creates thoughtful objects, installations, and experiences through contextual inquiry, visual data explorations, and UI/UX design. Contact her at malikaxdesign@gmail.com.

NOAH DEUTSCH is currently with Harvard University, Cambridge, MA, 02138, USA. He designs and builds software tools to address complex science and engineering challenges, including novel interfaces for autonomous spacecraft command. Contact him at ndeustch@mde.harvard.edu.

BENJAMIN P. S. DONITZ is currently with the California Institute of Technology, CA, 91109, USA. He works with the Jet Propulsion Laboratory in the Project System Engineering and Formulation section. Since joining JPL in 2019, he has worked on several mission concepts and studies evaluating the feasibility of sending dedicated spacecraft to interstellar objects and long period comets. Contact him at Benjamin.P.Donitz@jpl.nasa.gov.

OSKAR ELEK is currently with the University of California, Santa Cruz, CA, 95064, USA. He is a cross-disciplinary scientist specializing in complex systems simulation, artificial life, data modeling and visualization. Contact him at oskar@elek.pub.

SCOTT DAVIDOFF is currently with the California Institute of Technology, CA, 91109, USA. He uses the methods of human-computer interaction, data science, and human-centered design to investigate big scientific research questions in astrobiology, and planetary and earth science, and to understand the robots and teams that collaborate around exploration. He has invented problem solving techniques practiced at the world's leading companies and taught at its leading universities. Contact him at sd@scottdavidoff.com.

Computing in Science & Engineering

The computational and data-centric problems faced by scientists and engineers transcend disciplines. There is a need to share knowledge of algorithms, software, and architectures, and to transmit lessons-learned to a broad scientific audience. *Computing in Science & Engineering (CiSE)* is a cross-disciplinary, international publication that meets this need by presenting contributions of high interest and educational value from a variety of fields, including physics, biology, chemistry, and astronomy. *CiSE* emphasizes innovative applications in cutting-edge techniques. *CiSE* publishes peer-reviewed research articles, as well as departments spanning news and analyses, topical reviews, tutorials, case studies, and more.

Read *CiSE* today! www.computer.org/cise

