

Romie: A Domain-Independent Tool for Computer-Aided Robust Operations Management

Michael Saint-Guillain^a, Jonas Gibaszek^a, Tiago S. Vaquero^b, Steve A. Chien^b

^a Université catholique de Louvain, Place de l'Université 1, 1348 Ottignies-Louvain-la-Neuve, Belgium

^b Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109 USA
michael.saint@uclouvain.be, {tiago.stegun.vaquero, steve.a.chien}@jpl.nasa.gov

Abstract

Romie is a decision support tool based on AI's latest advances in the domain of robust scheduling. Unlike all its predecessors, the tool allows to (i) *visually model* the operational problem and context entirely (ii) optimize to find near-optimal schedules while *taking uncertainty into account* and deals with (iii) a combination of various *key performance indicators* (KPIs). It comes with a web user interface. Part or all of the modelled activities may be associated to random variables describing their stochastic durations, in order to produce schedules that are *robust* w.r.t. temporal uncertainty. Hence, depending on the pursued KPIs, the schedules maximize a combination of the following terms: the probability of satisfying the problem constraints, the expected return/efficiency, the expected outcome quality, and even the *operators' wellness* by minimizing its expected extra-hours. Initially developed for spatial exploration and demonstration in the context of Mars analog missions (i.e. missions on Earth that simulate condition and aspects of Mars missions), this *versatile* tool is here applied to operations management in both biotechnology manufacturing and robots parametrization in a cave exploration context.

Introduction

Project management realizes about 30% of the world gross product (Turner et al. 2010). However, most of existing studies have solely been done in machine scheduling environments (Herroelen and Leus 2005). How to deal with processing time uncertainty when facing a larger, complex, scheduling problem which possibly involves multiple human and/or operators, unknown probability distributions, hard deadlines and exotic constraints? In this paper, we introduce a visual tool, Romie, for computer-aided operations scheduling under uncertainty, and show how it has been successfully applied to three very different case studies of real world human operations management: 1) the *UCL to Mars 2018*¹ analog mission that took place at the Mars Research Desert Station (Utah), 2) the modelling and scheduling of a manufacturing project in an real Belgian biotech company, and finally 3) the complex mission operations of the Jet Propulsion Laboratory (JPL) team in the DARPA Subterranean Challenge. The results obtained from our three case studies convey three very important messages: (a) Even for very complicated and

¹UCLouvain (Belgium) university. <https://marsuclouvain.be>

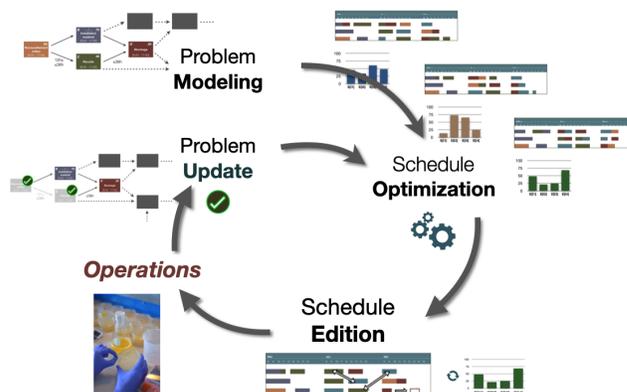


Figure 1: Principal functionalities of *Romie* tool.

various different operational contexts, a common modelling framework exists, being user friendly, visual, and rigorous at the same time; (b) Even for real sized problems, computer-optimized solutions outperform the schedules hand-crafted by field experts in general, and serve as a strong basis for decision making, as the deciders can always adapt and reuse these depending on external factors; (c) Schedules obtained while taking uncertainty into account systematically outperform that obtained from deterministic assumptions in terms of reliability and expected KPIs, while preserving most of the solutions quality; the latter result remains valid even when provided very bad representation of the uncertainty.

Description of the tool

Romie is a software tool aiming at supporting the decision makers in their operations management tasks. Romie uses combinatorial optimisation, in order to generate and optimize robust plannings for daily operations. The current key functionalities are depicted in Fig. 1:

- The user friendly, *visual modelling* of the problem at stake, in its own operational context: human and physical resources, operational constraints, key performance indicators (KPIs), execution uncertainties.
- *Robust scheduling*: the optimization engine takes the time uncertainty about each operation's duration into consideration, using modified-PERT distributions, yielding sched-



Figure 2: Left: the Mars Desert Research Station in Utah. Right: extra-vehicular field operations.

ules with high probability of success.

- The schedules are optimised while pursuing (a combination of) *various KPIs*: success probability, expected cost, expected quality, and even operators wellness.

In a research domain in constant evolution, *Romie* integrates state-of-the-art advances in robust scheduling under uncertainty (Saint-Guillain et al. 2021). Future versions will enable online monitoring of the operations, keeping the schedule and the underlying model consistent with the current state of the system, allowing the user to adapt and reoptimize future decisions based on past outcomes.

Optimization engine. A local search (LS) based approach, exploiting well-known sequence neighborhood operators (relocate, 2-opt, swap, ...) and a simulated annealing meta-heuristic, is used to explore the solution space. The LS algorithm uses sample average approximation (Kleywegt, Shapiro, and Homem-de Mello 2002) to evaluate the expected quality of a solution under time uncertainty. Each new solution is compared with the incumbent one in terms of its relative gain in each element of the ordered set of KPIs.

Timeline of Case Studies

This study presents a novel scheduling tool through three different case studies. Each case study happened sequentially, following and further validating different stages of the tool's development. Section 1 describes the *UCL to Mars 2018* mission case study, which marked the very first stage of the technology, assessing its ability to deal with complex projects made of simple operational tasks, constraints and resources. Section 2 describes a case study which took place from in 2019 and 2020, in collaboration with a Belgian biotech company. It permitted to extend our scheduling formalism and technology and reach a higher stage of applicability, in the complex real world industrial context. At the time these first two studies were conducted, the tool only consisted in a theoretical background, a modelling formalism and a versatile scheduling engine. Namely, there was no user interface (UI). The technology was validated, but not the ability of the end user to control and use it. Section 3 describes the third case study, in which human operators and robots from JPL collaborate in the final circuit of the DARPA Subterranean Challenge, scheduled for September 2021. *Romie* has been provided a brand new UI prototype, allowing (for the first time) an end-user to model the problem using a user-friendly visual interface, run optimization processes and visualize optimized schedules.

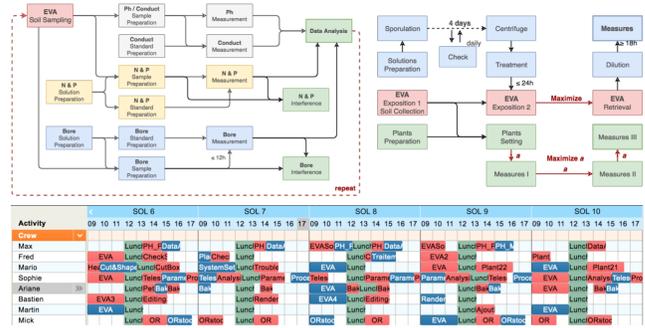


Figure 3: Top: visual models of two (out of seven) experiments conducted during the *UCL to Mars 2018* analog mission at MDRS. Top left: a soil analysis project in chemistry. Top right: a project mixing biology and botanic. Bottom: part of an optimized schedule followed the analog mission.

1 Robust Operations Management on Mars

The development of the scheduling tool started with the *UCL to Mars 2018* project. Unlike most scheduling problems, operations in a space mission must be planned days ahead. Complex decision chains and communication delays prevent schedules from being arbitrarily modified, hence online reoptimization approaches are usually not appropriate. The problem of scheduling a set of operations in a constrained context such as the *Mars Desert Research Station* (MDRS, Fig. 2) is not trivial, even in its classical deterministic version. It should be seen as a generalization of the well-known NP-complete *job-shop scheduling problem* (Lenstra and Kan 1979), which has the reputation of being one of the most computationally demanding (Applegate and Cook 1991). (Hall and Magazine 1994) reinforces the importance of mission planning, as 25% of the budget of a space mission may be spent in making these decisions beforehand, citing the Voyager 2 space probe for which the development of the a priori schedule, involving around 175 experiments, required 30 people during six months. Nowadays, hardware and techniques have evolved. It is likely that a super-equipped (*i.e.* with a brand new laptop) human brain suffices in that specific case. Yet, the problems and requirements have evolved too. Instead of the single machine Voyager 2, space missions have to deal with teams of astronauts.

Scheduling a Space Mission under Uncertainty. The purpose of the *UCL to Mars 2018* analog mission being to simulate intensive scientific activities in a extra-planetary context, the mission was organized based on 7 different research projects to be conducted at the MDRS, from various fields including biology, particle physics, medicine, engineering, botanic, chemistry and finally AI. In total, more than 230 tasks were involved by the seven research projects, with at least as many constraints. In fact, the modelling each research project merged within a global problem was inevitable, since all the activities at MDRS depend on the same limited and shared set of resources. Some projects required *extra vehicular activities* (EVAs), which for security

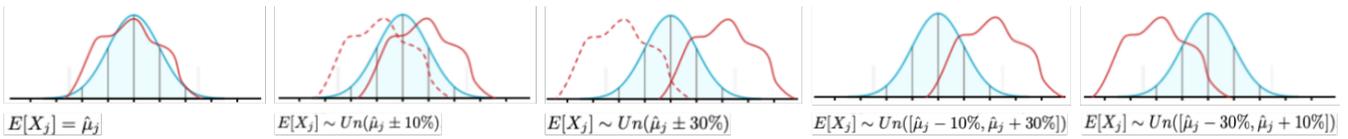


Figure 4: Varying the quality of the probability distributions, leading to five different experimental assumptions. Blue: estimator distribution, used at optimization stage. Red, both plain and dotted: real hidden distributions, revealed at execution stage.

reasons require at least three participants amongst the operators. EVAs usually take half a day, should be planned and approved days ahead and happen at most once a day. All operators then had their schedule linked to each others, even concerning research projects that do not require EVAs. Fig. 3 shows the modelling of two such research projects, together with an excerpt of optimized schedule. At the MDRS however, computing an optimal schedule becomes significantly less attractive as problem data, such as the manipulation time of experiments, are different from their predicted values. In a constrained environment with shared resources and devices, such deviations can propagate to the remaining operations, eventually leading to global infeasibility. In (Saint-Guillain 2019), we investigated based on the real case study of a Mars analog mission, the impact of stochastic robust modeling against a classical deterministic approach on the reliability of a priori mission planning.

Uncertainty and Performance indicators. The main purpose of (Saint-Guillain 2019) was to compare both deterministic and robust stochastic approaches to the problem of scheduling a set of scientific tasks under processing time uncertainty, in the operational context of a Martian planetary habitat. We empirically showed that taking uncertainty into account while optimizing the schedules allows significant gains on average when applied on real instances involving the constraints faced and objectives pursued during a two-week Mars analog mission. The objectives were both optimizing the mission’s success probability, namely the *robustness*, in terms of meeting the operational constraints and deadlines, and maximizing the total scientific outcome, as a linear combination of specific metrics, designed according to the scientific objectives of each of the seven research projects. Therefore, the optimization engine was computing a near-optimal schedule when optimizing the robustness KPI first, and expected scientific outcome second.

Uncertainty on the Uncertainty. The computational experiments did also take into account uncertainty on the stochastic knowledge itself. In fact, accurate estimations of the probability distributions require a significant amount of observations, which is often impossible in practice. It is then usually both necessary and realistic to consider the real distributions as unknown (or hidden). Consequently, in (Saint-Guillain 2019) the same computational experiments were reiterated under five different assumptions about the estimator’s quality. The concept is illustrated in Fig. 4: five hypothetical couples of both estimator and real distributions are drawn. The estimator distribution is simply the distribution

used to describe the duration of an activity. It represents the current knowledge one has about the activity’s uncertainty. The estimator distributions are the only information available at optimization phase. Real (hidden) distributions are only revealed when the computed schedules, therefore optimized considering blue distributions, were executed in a simulation. The shapes of the real distributions were always randomly generated, while controlling the resulting mean. In the first case, the estimators (which used to be normal distributions) were of very good quality, since their means always coincide with that of the real distributions (even if the shapes differ). In the third case, these were of bad quality, since each estimator mean may be up to 30% away from that of the associated real distribution. Finally, the fourth (*resp.* fifth) case stands for the situation in which all durations are underestimated (*resp.* overestimated) in general. The results revealed that, even when using very bad approximations of the probability distributions, the computed solutions significantly outperform those obtained from a classical deterministic formulation, while preserving most of the solution’s quality.

Experiments and Results. The aforementioned paper gives extensive details on both the experimental plan and results. In a nutshell, we observed that as the accuracy of the probability distributions that describe the project task durations varies from really accurate (no under/over estimation on average) to very bad (30% under/over estimations on average), the proportion of the simulations in which the schedules computed based on a classical deterministic model lead to a successful execution varies from 5% to 8% only. Using our probabilistic model, these success rates increase significantly, between 70% and 90%. Two additional assumptions were tested, in which the durations were either systematically underestimated, or systematically overestimated. The first case is naturally catastrophic, leading to success of approximately 0.1% for the deterministic schedules, and 22% thanks to our stochastic model. The second case is really interesting, as it describes a very common behaviour of managers who have to face time uncertainty, which consists in systematically considering a duration larger than what they believe the tasks is likely to last. In this particular context, our simulations showed that our stochastic model increases the schedules’ robustness from 34% to 95%.

Benefits and Price of Robustness. The experimental results clearly indicate the benefits of our probabilistic model over a deterministic one, as even in the ideal case of all durations being overestimated on average, the resulting schedules reveal three times more reliable (succeed in 95% of the

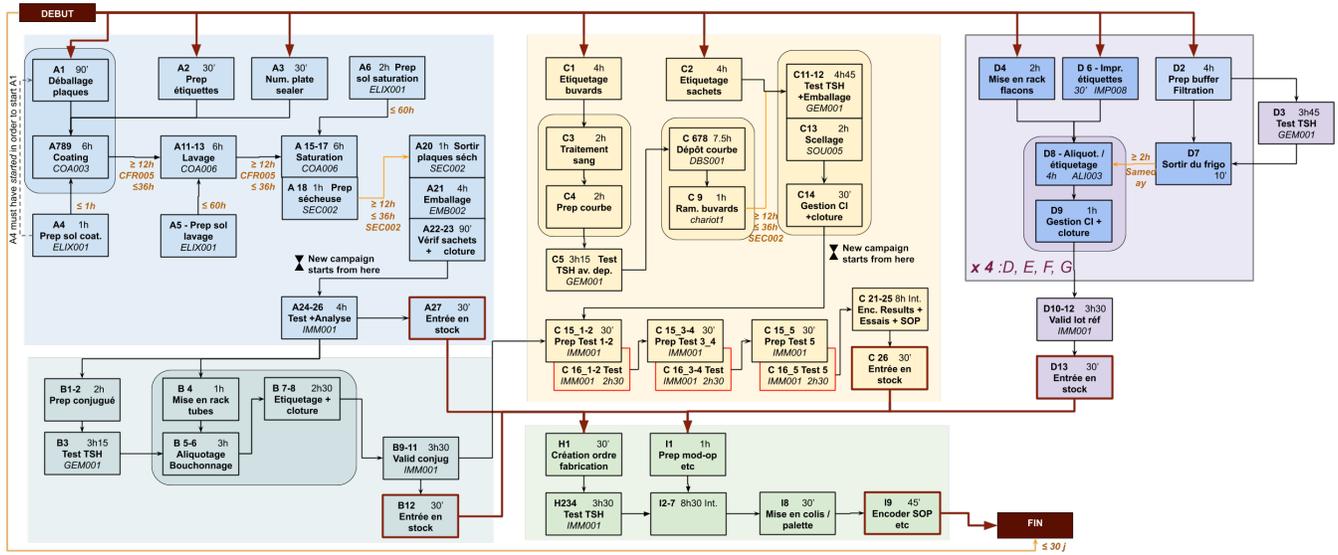


Figure 5: A visual model that fully describes the tasks and constraints involved in one production campaign of the biotech company’s most popular product.

cases). Independently of the estimators accuracy, the relative difference in terms of the scientific outcome KPIs, between schedules produces from both models, is of only 7% on average. In other words, even all durations being overestimated and in our MDRS context, a deterministic schedule yields 7% more science on average, whereas it has at 60% chances of failing the mission. This is the value of perfect information over uncertainty. On the other hand, a success rate of 95% can be reached by sacrificing 7% of the outcome. This is the benefits, and price, of robustness.

2 Biotech Manufacturing

As the day humans will live on Mars is still far ahead, we wanted to extend our tool as well as the underlying technology to tackle significantly different operational context, others than human space missions. A Belgian company, specialized in biotechnology product manufacturing, accepted to collaborate on the concrete project of modelling the scheduling problem involved in the manufacturing of one of their most popular products, and eventually solving this scheduling problem, at different scales.

Modelling in the Industrial Real World Contexts. Fig. 5 shows how their production problem was modelled, using the exact same *visual formalism* than that used in the context of the *UCL to Mars 2018* mission. More precisely, the diagram depicted only represents one single production campaign, involving around 85 tasks. However, for efficiency reasons the company would usually run up to three production campaigns in parallel, whereas the operational human and physical resources remain fixed.

As the operational context in a biotech company significantly differs from that of a space station, various additional exotic constraints had to be added to the formalism,

and the other scheduling optimization technology extended consequently. For instance, running several production campaigns in parallel involve specific constraints, stating how these may overlap or not. Naturally, we had to cope with complex operator worktime management, such as weekends, days off, part times, *etc.*. In particular, the concept of extra working time, whenever an operator must remain on site later than normally accepted, revealed of major importance. Finally, as the production process requires many tasks to be double checked, constraints of the form “task A cannot be executed by the same operator than task B” were mandatory too, and special constraints related to physical resources had to be designed as well. These are just examples amongst the large number of additions that extended the initial technology. Fig. 6 shows an example of solution for three production campaigns in parallel.

Problem Modelling made Simple. The modelling formalism used to communicate with our collaborators from the biotech company, namely the language used to describe their scheduling problem, is based on simple diagrams such as depicted in Fig. 5. In practice, once the problem visually modelled, we translated the diagrams in a mathematical language accepted by our scheduling engine. It is worth noting that the time needed by our collaborators, being not scheduling experts, to master the proposed visual modelling formalism (and autonomously draw part of the models) revealed to be of five to ten one-hour meetings only.

Optimizing Wellness: Stress Aversion. Cost-based KPIs, such as minimizing total production makespan, or quality-based KPIs such as the metric of total scientific outcome considered in the *UCL to Mars 2018* case study, are classical objectives to be pursued. In a human context however,

		% Success			Extra-Hours				Makespan (days)			
		Exact	Under	Over	Plan	Exact	Under	Over	Plan	Exact	Under	Over
1C Det.	Min. makespan	6.6	0.9	17.1	9.9	9.7	11.3	10.0	15.1	15.7	16.6	15.4
	Min. extra-hours	60.9	37.5	40.2	0.0	4.6	5.9	4.3	17.2	19.8	20.7	18.5
1C Stoch.	Min. makespan	98.4	83.7	98.0	1.4	4.1	6.1	3.1	17.2	17.2	17.3	17.2
	Min. extra-hours	99.9	99.6	100.0	1.0	1.9	2.7	1.0	26.9	27.1	27.4	27.0
3C Det.	Min. makespan	0.5	0.0	2.5	16.6	29.1	34.6	26.7	38.1	41.9	42.5	40.1
	Min. extra-hours	14.2	2.2	12.5	1.0	15.7	23.6	11.3	45.5	46.4	46.8	45.9
3C Stoch.	Min. makespan	96.2	61.1	96.0	6.2	13.0	20.0	10.3	43.2	43.2	43.3	43.2
	Min. extra-hours	99.4	97.7	98.7	3.0	6.4	9.9	3.8	58.9	58.9	58.9	58.9

Table 1: Biotech manufacturing case study, involving one (1C) and three (3C) production campaigns, optimizing under either deterministic (Det.) or stochastic model (Stoch.). We consider 3 different assumptions about the stochastic knowledge: exact mean, 10% under-estimations and 10% over-estimations, as described in Section 1: first, fourth and fifth in Fig. 4, this time with $E[X_j] \sim U(\mu_j - 3\%, \mu_j + 10\%)$ and $E[X_j] \sim U(\mu_j - 10\%, \mu_j + 3\%)$. Statistics include the percentages of simulations in which the optimized schedules (% Success) are found to respect all the problem constraints when executed under the “hidden” uncertainty (red in Fig. 4). The *Plan* columns indicate the KPI values as predicted by the *a priori* schedule, as opposed with the average KPI values observed during simulations (other columns). *Minimize makespan*: optimization done while minimizing the makespan first, extra-hours second. *Minimize extra-hours*: extra-hours first, makespan second.

to rapidly map, navigate, and search underground environments including natural cave networks, tunnel systems, and urban underground infrastructure for particular objects of interest (e.g. mannequin survivors, backpacks, cell phones, helmets, etc.), called *artifacts*. Subterranean environments pose significant challenges for manned and unmanned operations due to limited communication and situational awareness. In the SubT competition in particular, only a single human operator is allowed to interact with the robotic team. CoSTAR’s robotic team consists of more than four robots with wheeled, legged, and flying mobility.

Operating multiple robots with different capabilities in kilometer-long underground environments can go beyond the cognitive capacity of a single human supervisor – even with advanced autonomy in place. SubT operations may involve cognitively demanding tasks such as monitoring 3D mapping of the environment and localization accuracy, establishing communication links between robots, assessing location and health of all robots, and submitting detected artifacts within the allotted competition time. In order to facilitate operations during the SubT competition, CoSTAR team has developed the Copilot MIKE (Kaufmann et al. 2021), an autonomous assistant for human-in-the-loop multi-robot operations. During complex and potentially stressful exploration missions, MIKE helps by planning operation tasks related to setting up and commanding the robots, while maintaining a bearable workload and high situational awareness.

In this experiment we study the use of *Romie* to i) model the tasks and constraints that are required during setup time and during competition, and to ii) support the operator by scheduling the tasks in a way that maximizes the robustness. This case study has principally differs from the two previous ones by the scale of its time horizon to manage. Whereas the schedules at MDRS, as well as in the context of biotech production, typically involve several weeks, the DARPA SubT lasts only one hour. In particular, the problem at stake consists here at both 1) setting and getting all the exploration

robots for starting the mission (setup time) and 2) deploying the team of robots to explore the target underground environment. The scheduling problem involves setup tasks having nominal durations ranging between 10 and 120 seconds.

Problem Modelling. Figure 7 shows a graphical drawing of the problem structure. This diagram, as well as those shown in Figures 3 and 5, have been drawn using *Google-Docs Drawings*. The model involves two principal groups of activities: the *Base group* (in blue) and the *Robot group* (in yellow). The Base group includes unique activities that are common to all robots. The Robot group represents all the activities specific to one robot. If we have 7 robots, there are 7 duplicates of each activity from the Robot group. The main challenge results in the fact that everything must be scheduled in order to fulfill the time constraints, depicted in red. For instance, the *Start Robot Logging* activity must append between 9:28am and 9:30am, for every robot. Finally, the human operator is assumed to be able to carry on up to 4 activities at the same time, whereas some activities (denoted by * in Fig. 7) can be entrusted to MIKE, which could follow up to 5 activities in parallel. It is important to note that the only key performance index (KPI) here is the probability of success, that is, the probability that the schedule actually respects all the temporal constraints shown in red in Fig. 7. In fact, contrary to previous case studies, the problem at stake here is no longer an optimization problem but a constraint satisfaction problem, in the sense that there is no other goal to be pursued than satisfying these constraints.

The Modelling User Interface. Our tool now integrates a visual modelling user interface, depicted in Figure 8. The UI allows to describe the scheduling problem structure, in terms of temporal and operational constraints. It also allows to input constraints and activity properties, which were not present in the hand-drawn diagram of Fig. 7. In particular, the model allows for modified-PERT distributions, a

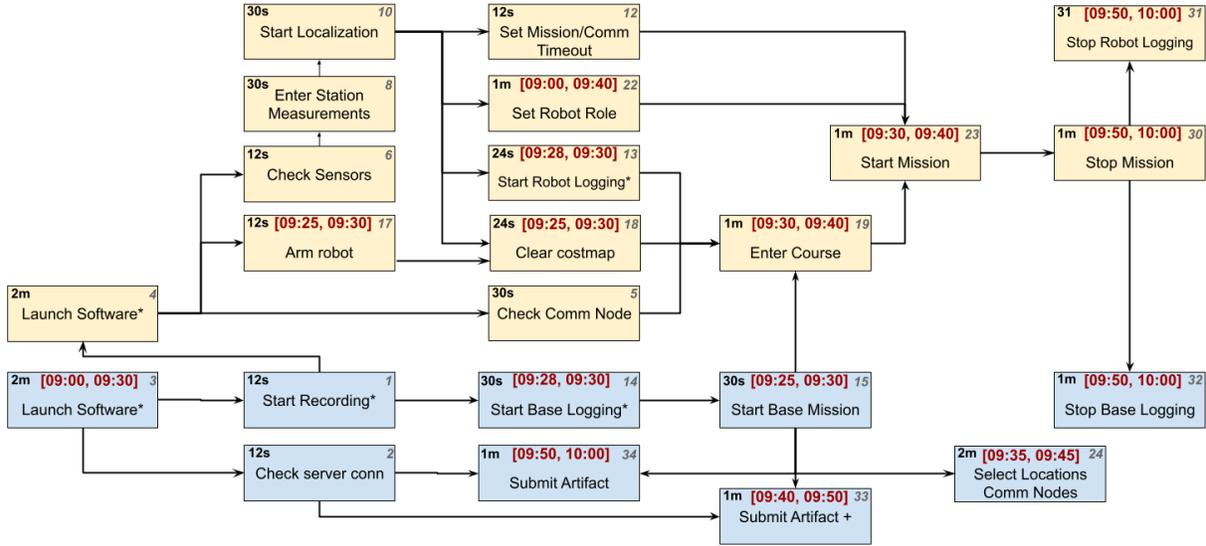


Figure 7: The DARPA SubT problem structure, as designed by hand using a classical diagram editor.

probability distribution widely used in risk analysis (Kamburowski 1997), which has the advantage of enabling asymmetric bounded probability distributions.

Computing Robust Schedules. Our scheduling engine exploits parallelization paradigm to achieve reasonable computation times *Romie* is capable of exploiting parallel computing. On a 60 cores cluster, finding feasible solutions (*i.e.* schedules that fulfil the temporal constraints listed in Fig. 7) takes less than 15 seconds only for problems with up to 15 robots. Up to 18 robots, it takes less than a minute to find its first feasible solution. Once an initial solution has been found to be deterministic feasible, meaning that it is feasible when considering all its durations as deterministic, the engines switch on the probabilistic optimization mode and pursue its duty while optimizing with respect to the expected KPIs: here the probability of success only. No deterministic feasible solution was found for 19 robots or more (in less than five minutes). An optimized schedule is depicted in Fig. 9, using the tool’s integrated visualization interface.

Still exploiting 60 cores, for problems involving up to 12 robots *Romie* finds *optimal* solutions in less than 5 seconds. Of course, the nature of the scheduling engine embedded in *Romie*, a local search based solution framework, naturally prevents from providing any optimality proof in general. For the DARPA SubT challenge however, the only optimized KPI being the success probability, solutions with 1.0 success probability are necessarily optimal. In our case, *Romie* finds such solutions for instances involving up to 15 robots, in which 20 seconds only are required to find a first deterministic feasible solution, and optimal solutions are found within approximately 2 minutes.

Average Performances of the Computed Schedules. As usual, the optimized robust schedules are compared to solutions optimized based on deterministic assumption, there-

Bots	1	4	6	10	15	16	17	18
Exa	91.0	82.8	67.8	31.9	0.1	0.0	0.0	0.0
Und	90.0	80.8	58.5	17.4	0.0	0.0	0.0	0.0
Ovr	93.3	88.1	81.4	52.7	9.6	7.0	4.8	0.2
Exa	100	100	100	100	100	95.6	61.5	0.3
Und	100	100	100	100	93.0	40.7	1.3	0.0
Ovr	100	100	100	100	100	100	97.9	56.4

Table 2: DARPA SubT challenge case study, involving up to 18 robots. We consider three different assumptions about the stochastic knowledge: exact mean (Exa), [-3%,+10%] under-estimations (Und) and [-10%,+3%] over-estimations (Ovr), as described for Table 2. First rows give the average percentage success, when schedules are optimized using a deterministic model. The second set of rows give average results when *Romie* uses the stochastic model.

fore measuring the average gain at optimizing based on a probabilistic model against a classical, deterministic one. The average results are provided in Table 2. Solutions optimized under the stochastic model significantly outperform that of the deterministic one, in terms of reliability (*i.e.* robustness), that is, the average percentage of simulation success under the three different assumptions made on the quality of the stochastic knowledge (*i.e.* accuracy of the chosen PERT distribution parameters for each activity duration).

4 Conclusions and Future work

In this paper we presented *Romie*, a state-of-the-art robust scheduling tool, based on the principle of optimizing under time uncertainty. We described three very different application cases that were handled with our tool, showing the versatile aspect of *Romie*, which allows to model and solve scheduling problem despite the different operational contexts. The more recent case study shows that end users are

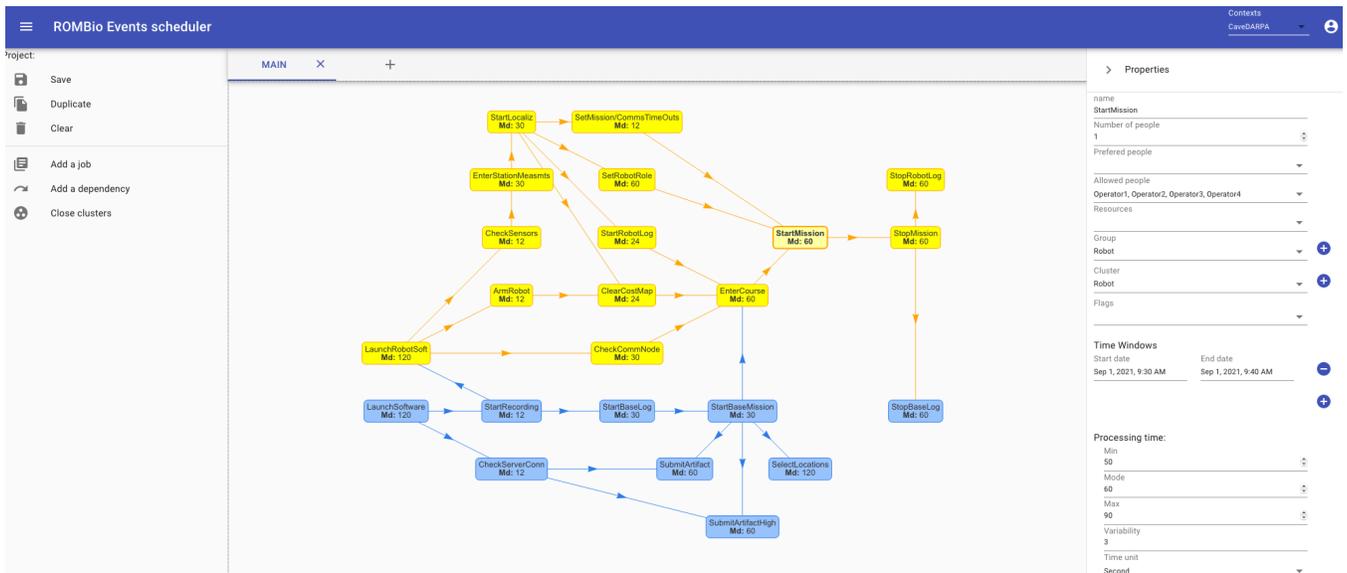


Figure 8: Modelling using the *Romie* tool’s user interface. Amongst the different parameters that define an activity: allowed operators, used resources, execution time windows, parameters defining the duration’s probability distribution (PERT).

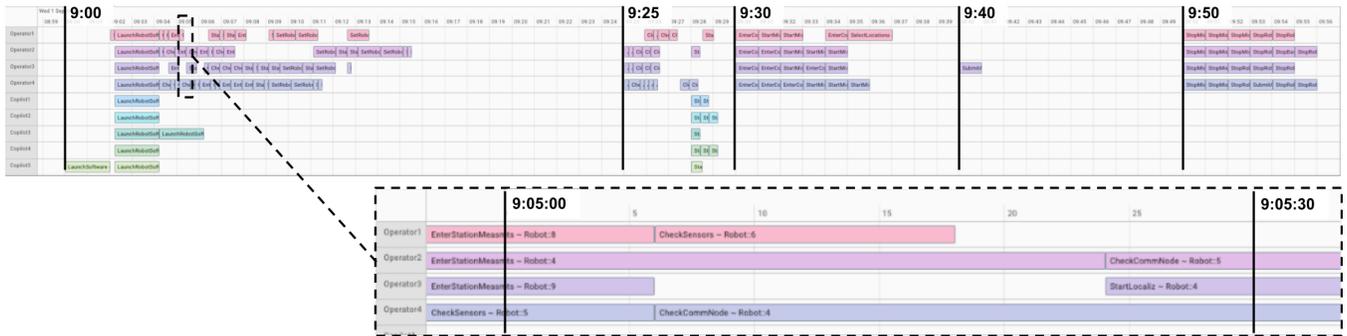


Figure 9: An example schedule solving the DARPA Subterranean robot problem, showed in the *Romie* tool’s schedule visualization interface. Top: 1 robot problem. Middle: 10 robots problem. Each line corresponds to an operator timeline (4 humans, 5 copilots). The interface allows the user to zoom at will, has showed in the bottom of the figure.

actually able to visually describe the scheduling problem at stake, and further visualize optimized solutions. Once again, the benefits of using a probabilistic modelling approach, taking the time uncertainty of activities durations into account at optimization stage, are clearly confirmed by the empirical average gains compared to a classical deterministic approach. Our approach also allows to significantly decrease extra hours and deviations from a priori decisions, hence reducing the operators’ stress load in manufacturing context.

Future development directions. Our tool currently integrates the key functionalities for *a priori* robust scheduling under uncertainty. The next logical step will naturally be to integrate online management functionalities, namely *monitoring* and *reoptimization*. Online monitoring of the operations aims at updating the schedule, as well as the underlying model, in order to keep them consistent with the current state of the operations. It eventually allows the possibility for

adapting and reoptimizing future decisions, in light of past decisions and outcomes. These additional features will be useful for the scheduling of future research and operational projects. Applications to biotech online manufacturing problems are planned in the near future, as well as the next circuit of the DARPA Subterranean Challenge (September 2021).

Acknowledgements. We thank the *UCL to Mars 2018* crew: Bastien Baix, Frédéric Peyrusson, Martin Roumain, Ariane Sablon, Mario Sundic, Sophie Wuyckens and Maximilien Richald. We thank our collaborators from the Belgian biotech company *Zentech*: Isabelle Sansdrap, Nursel Yilmaz. We thank the Mars Society, Shannon Rupert and Mission Support personnel, for creating the conditions for realistic Mars analog sojourns. The research was partially carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

References

- Applegate, D.; and Cook, W. 1991. A computational study of the job-shop scheduling problem. *ORSA Journal on computing* 3(2): 149–156.
- Hall, N. G.; and Magazine, M. J. 1994. Maximizing the value of a space mission. *European journal of operational research* 78(2): 224–241.
- Herroelen, W.; and Leus, R. 2005. Project scheduling under uncertainty: Survey and research potentials. *European journal of operational research* 165(2): 289–306.
- Kamburowski, J. 1997. New validations of PERT times. *Omega* 25(3): 323–328.
- Kaufmann, M.; Vaquero, T.; Correa, G.; Otsu, K.; Ginting, M.; Beltrame, G.; and Agha-Mohammadi, A.-A. 2021. Copilot MIKE: An Autonomous Assistant for Multi-Robot Operations in Cave Exploration. In *IEEE Aerospace Conference*.
- Kleywegt, A. J.; Shapiro, A.; and Homem-de Mello, T. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2): 479–502.
- Lenstra, J. K.; and Kan, A. R. 1979. Computational complexity of discrete optimization problems. In *Annals of discrete mathematics*, volume 4, 121–140. Elsevier.
- Saint-Guillain, M. 2019. Robust operations management on mars. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, 368–376.
- Saint-Guillain, M.; Stegun Vaquero, T.; Chien, S.; and Abrahams, J. 2021. Probabilistic Temporal Networks with Ordinary Distributions: Theory, Robustness and Expected Utility. *Journal of Artificial Intelligence Research (under publication)*.
- Turner, R. J.; Huemann, M.; Anbari, F. T.; and Bredillet, C. N. 2010. *Perspectives on projects*. Routledge.