







Sampling Distributed Schedulers for Resilient Space Communication

Pedro R. D'Argenio^{1,2,3} , Juan A. Fraire^{1,2,3} , and Arnd Hartmanns⁴  

¹ CONICET, Córdoba, Argentina

² Saarland University, Saarbrücken, Germany

³ Universidad Nacional de Córdoba, Córdoba, Argentina

⁴ University of Twente, Enschede, The Netherlands

a.hartmanns@utwente.nl

Abstract. We consider routing in delay-tolerant networks like satellite constellations with known but intermittent contacts, random message loss, and resource-constrained nodes. Using a Markov decision process model, we seek a forwarding strategy that maximises the probability of delivering a message given a bound on the network-wide number of message copies. Standard probabilistic model checking would compute strategies that use global information, which are not implementable since nodes can only act on local data. In this paper, we propose notions of distributed schedulers and good-for-distributed-scheduling models to formally describe an implementable and practically desirable class of strategies. The schedulers consist of one sub-scheduler per node whose input is limited to local information; good models additionally render the ordering of independent steps irrelevant. We adapt the lightweight scheduler sampling technique in statistical model checking to work for distributed schedulers and evaluate the approach, implemented in the MODEST TOOLSET, on a realistic satellite constellation and contact plan.

1 Introduction

There is an increasing commercial and scientific interest in deploying large-scale satellite networks in low-Earth orbit (LEO) to collect and distribute information [4]. Real-time access to data is, however, only feasible when many satellites align to form a chain of links between a (remote) destination source or destination and a ground station. This vision favours mega-constellations; e.g. SpaceX's Starlink is composed of 12,000 satellites [26]. A different and more sustainable approach is to relax the real-time constraint and leverage the store-carry-and-forward principle where nodes store received messages for later forwarding to other nodes in the network, once a communication window—a contact—appears. This gives rise to a *delay-tolerant network* (DTN) [19]. Originally intended for

The authors are listed alphabetically. This work was supported by ANPCyT PICT-2017-3894 (RAFTSys), ERC grant 695614 (POWVER), NWO VENI grant no. 639.021.754, and SeCyT project 33620180100354CB (ARES).

© Springer Nature Switzerland AG 2020

R. Lee et al. (Eds.): NFM 2020, LNCS 12229, pp. 291–310, 2020.

https://doi.org/10.1007/978-3-030-55754-6_17

interplanetary networks [16], the DTN architecture has been identified as a disruptive approach for LEO constellations which allows for a better utilisation of communication opportunities [17]. In DTN, there is no upper bound on the propagation delay, and no expectation of continuous or bi-directional end-to-end connectivity. While DTN satellite constellations do not work for e.g. voice services, they can network Earth observation and high-latency data service missions. In particular, they can utilise slower inter-satellite links by advance forwarding [28], adapt transmission schedules to limited battery conditions [31], work with constrained spacecraft antennas and subsystem architectures [32], and tailor communication resources to fit mission traffic and routes [33,34]. Ultimately, the DTN approach enables sparse topologies of fewer and cheaper satellites [30].

A *contact* is the opportunity to establish a temporal communication link. As a DTN node’s network state information may be inaccurate or obsolete, traditional Internet routing schemes cannot be used. Space DTN routing approaches (both near-Earth and deep-space) thus seek to exploit the *a priori* knowledge of contacts: inter-satellite and satellite-to-ground contacts can be precomputed based on the orbital elements and communication parameters. The result is a *contact plan* [27]. We visualise a simple plan for four satellites—nodes N_1 through N_4 —in Fig. 1. A bent vertical arrow indicates a contact from the arrow’s origin to its target node. We abstract real time into discrete time slots T_1 through T_5 ; actual contact plans would show actual time intervals of varying durations (and potentially overlapping) with up to sub-second precision. Contact plans describe the expected network connectivity over time. They are the input for centralised or distributed DTN routing procedures. Existing solutions turn contact plans into e.g. time-expanded graphs [53] or contact graphs [5] on which routing calculations can be performed efficiently. Contact graphs have notably been validated by technological demonstrations in orbit [45,54].

In practice, the actual contacts may differ from the original plan due to failures or incomplete/inaccurate knowledge at the time the plan was computed. Space DTN in particular face fault-prone nodes [29], interference-sensitive communication links [52], and inaccurate orbit determination and station keeping procedures [39]. We thus need *uncertain contact plans* where contacts may fail for various reasons. Based on statistical data, we can annotate every contact with its success probability. In Fig. 1, we use probabilities p_1 through p_5 for illustration. Given an uncertain contact plan, we would then like to find a routing strategy that maximises the probability that the message is delivered to its destination. To increase that probability, we can allow *copies* of the message that




Fig. 1. Abstract uncertain contact plan

propagate along different paths. However, as typical DTN satellites have limited resources, we also want to bound the number of copies. Existing routing schemes only perform well under perfectly known or fully unknown contact plans [29, 49], leaving significant room for improvement for uncertain contact plan routing.

The routing problem in uncertain contact graphs with bounded copies matches very well with the modelling capabilities of Markov decision processes (MDP) [50]. They combine discrete probabilistic choices as in discrete-time Markov chains, which can represent contact failures, with nondeterministic choices as in Kripke structures [7], which can represent the routing options. Given an MDP model, we can use probabilistic model checking (PMC) [6] to determine the routing strategy (corresponding to the *scheduler* in PMC) with the highest probability for eventual message delivery. Raverta et al. [51] recently used this approach with optimisations that exploit the structure of the DTN routing models. However, PMC computes *global-information* schedulers, which take the local states of *all* nodes into account to make the optimal decision; they are thus *unimplementable*.

Example 1. In the contact plan of Fig. 1, assume we can send one message per slot. The highest-probability route from N_1 to N_4 is N_1 - N_2 - N_3 - N_4 (in slots T_1 , T_2 , and T_4 , with probability 0.405). The second-best is N_1 - N_3 - N_4 (in T_3 and T_4 , $p = 0.25$). Sending directly from N_1 to N_4 is least reliable. If N_1 starts with two copies of the message, then the first should be sent to N_2 in T_1 . In slot T_3 , N_1 can then either try to send the remaining copy to N_3 , or keep it for slot T_5 . We will show in Sect. 2 that the best choice for node N_1 computed by PMC is to send in T_3 iff node N_3 does not already have a copy of the message, i.e. if communication in slot T_1 or T_2 failed. In a space DTN, N_1 *cannot know this!*

PMC is thus not well suited for space DTN routing. The underlying problem of applying PMC to distributed systems was recognised almost 20 years ago [3, 21, 35], and led to the development of the notion of (*strongly*) *distributed schedulers* [35, 36] that only act on locally observable information. However, depending on the exact formalism and definition used, PMC for these schedulers is undecidable [35] in general, and NP-hard in the memoryless case [2, 36].


In this paper, we propose to use statistical model checking (SMC) [44, 56] with lightweight scheduler sampling (LSS) [48] in place of PMC to obtain routing strategies with a high probability for message delivery. Our contributions are (1) a modern and practical definition of distributed schedulers and models (in Sect. 3) appropriate for the space DTN setting that matches the compositional state-based modelling approach with undirected synchronisation common to today’s probabilistic modelling languages [9, 14, 40, 46] and tools [24, 41, 43, 46]; (2) an adaptation and implementation of SMC with LSS for distributed schedulers (in Sect. 4), and (3) a modelling pattern and SMC-based analysis toolchain for routing in space DTN with uncertain contact plans (in Sect. 5). We start (in Sect. 2) with a simplified but complete definition of the compositional MDP formalism that underpins our approach, introducing a detailed model for Fig. 1 along the way. We end (in Sect. 5) with an experimental evaluation of our new technique on a realistic satellite constellation model and contact plan.

2 Scheduling in Markov Decision Processes

Preliminaries. \mathbb{Z} is the set of integer numbers. We write $[a, b]$ for the real interval $\{x \in \mathbb{R} \mid a \leq x \leq b\}$. Given a set S , its powerset is 2^S . A probability distribution over S is a function $\mu: S \rightarrow [0, 1]$ with countable support $spt(\mu) \stackrel{\text{def}}{=} \{s \in S \mid \mu(s) > 0\}$ and $\sum_{s \in spt(\mu)} \mu(s) = 1$. $Dist(S)$ is the set of all probability distributions over S . We write $\{x_1 \mapsto y_1, \dots\}$ for the function that maps each x_i to y_i , and if necessary implicitly maps to 0 all other x . Thus we can e.g. write $\{s \mapsto 1\}$ for the *Dirac distribution* that assigns probability 1 to s . For a tuple $t = \langle y_1, \dots, y_n \rangle$, $t[i] \stackrel{\text{def}}{=} y_i$ is its i -th component.

Definition 1. A Markov decision process (MDP) is a tuple $M = \langle S, s_I, A, T \rangle$ where S is a finite set of states with initial state $s_I \in S$, A is a finite set of actions, and $T: S \rightarrow 2^{A \times Dist(S)}$ is the transition function with $T(s) \neq \emptyset$ for all $s \in S$. The set of all transitions is $Tr(M) \stackrel{\text{def}}{=} \cup_{s \in S} T(s)$; it must be finite.

We also write $s \xrightarrow{a}_T \mu$ for $\langle a, \mu \rangle \in T(s)$, and may omit the $_T$ subscript. An element of $spt(\mu)$ is a *branch* of transition $s \xrightarrow{a}_T \mu$. To leave a state, we first choose a transition, then select the next state probabilistically among its branches. An MDP with $\forall s: |T(s)| = 1$ is a *discrete-time Markov chain* (DTMC). We draw MDP as shown above on the right: this MDP has three states



with $s_I = s_0$, four transitions, and five branches. For transitions with a single branch, we omit the dot and probability 1. We have $s_0 \xrightarrow{a} \{s_1 \mapsto \frac{1}{3}, s_2 \mapsto \frac{2}{3}\}$.

Modelling with MDP directly is cumbersome; we instead use a higher-level modelling language like MODEST [9, 40] that extends MDP with discrete variables and parallel composition. Given a set of (integer-valued) *variables* X , let $Val_X \stackrel{\text{def}}{=} X \rightarrow \mathbb{Z}$ be their *valuations*. Let Bxp_X and Ixp_X contain all Boolean and integer *expressions* over the variables in X , respectively. We omit $_X$ subscripts if clear from the context. For $e \in Bxp$ ($e \in Ixp$), let $v(e) \in \{true, false\}$ ($v(e) \in \mathbb{Z}$) be the value of e in $v \in Val$. Finally, let $Upd \stackrel{\text{def}}{=} X \mapsto Ixp$ be the set of *updates* that map each variable to an expression determining the value assigned to it.

Definition 2. An MDP with variables (VMDP) is a tuple $M = \langle Loc, \ell_I, A, X, x_I, E \rangle$ where Loc is a set of locations with initial location $\ell_I \in Loc$, A is a set of actions, X is a set of variables with initial values given by $x_I \in Val$, and $E: Loc \rightarrow 2^{Bxp \times A \times Dist(Upd \times Loc)}$ is the edge function. All sets must be finite.

We write $s \xrightarrow{g,a}_E \nu$ for $\langle g, a, \nu \rangle \in E(\ell)$, and may omit the $_E$ subscript. An edge in a VMDP has a *guard* g that determines whether the edge is *enabled*. A branch of an edge carries an update u that changes the variables’ values. Formally:




Fig. 2. Four VMDP modelling the nodes of the example contact plan

Definition 3. Given a VMDP $M = \langle Loc, \ell_I, A, X, x_I, E \rangle$, its semantics is the MDP $\llbracket M \rrbracket \stackrel{\text{def}}{=} \langle Loc \times Val, \langle \ell_I, x_I \rangle, A, T \rangle$ with T the smallest function satisfying

$$\ell \xrightarrow{g,a} E \nu \wedge v(g)$$

$$\langle \ell, v \rangle \xrightarrow{a} T \{ \langle \ell', v' \rangle \mapsto \sum_{\{u | u \in Upd \wedge v' = \{x \mapsto v(u(x))\}\}} \nu(\langle u, \ell' \rangle) \mid \ell' \in Loc, v' \in Val \}$$

We must restrict to VMDP whose semantics is finite and deadlock-free.

Example 2. Figure 2 shows four VMDP N_1 through N_4 that model the nodes of Fig. 1. Every node has a variable c_i to track the number of message copies it owns. We write $x \leftarrow e$ for the mapping of variable x to value or expression e . In every slot where a node N_i can send, it has a choice between two transitions labelled nop_i (do not send) and snd_i (send one copy: decrement c_i , set d to 1). In a slot T_j where N_i can receive, it always tries to do so via action rcv ; this succeeds with probability p_j as given in Fig. 1. If the sender decided not to send, then a successful receive has no effect on c_i because d is zero. The *parallel composition* of these four VMDP models the entire contact plan, with the nodes synchronising on shared action rcv and exchanging data via shared variable d .

Definition 4. Given two VMDP $M_i = \langle Loc_i, \ell_{I_i}, A_i, X_i, x_{I_i}, E_i \rangle$, $i \in \{1, 2\}$, a finite set A of actions, and a synchronisation relation

$$\text{sync} \subseteq (A_1 \uplus \{\perp\}) \times (A_2 \uplus \{\perp\}) \times A,$$

their parallel composition is

$$M_1 \parallel_{\text{sync}} M_2 \stackrel{\text{def}}{=} \langle Loc_1 \times Loc_2, \langle \ell_{I_1}, \ell_{I_2} \rangle, A, X_1 \cup X_2, x_{I_1} \cup x_{I_2}, E \rangle$$

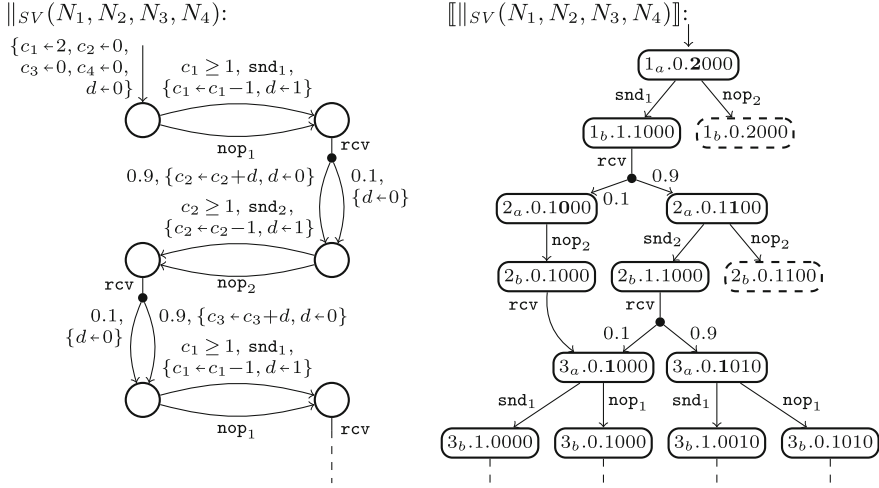


Fig. 3. Excerpt of the network of the node VMDP (left) and its semantics (right)

where E is the smallest function that satisfies the inference rules

$$\frac{l_1 \xrightarrow{g_1, a_1} E_1 \nu_1 \quad \langle a_1, \perp, a \rangle \in \text{sync}}{\langle \ell_1, \ell_2 \rangle \xrightarrow{g_1, a} E \{ \langle \langle \ell'_1, \ell_2 \rangle, u_1 \rangle \mapsto \nu_1(\langle \ell'_1, u_1 \rangle) \mid \langle \ell'_1, u_1 \rangle \in \text{spt}(\nu_1) \}} \quad (\text{ind}_1),$$

$$\frac{l_1 \xrightarrow{g_1, a_1} E_1 \nu_1 \quad l_2 \xrightarrow{g_2, a_2} E_2 \nu_2 \quad \langle a_1, a_2, a \rangle \in \text{sync}}{\langle \ell_1, \ell_2 \rangle \xrightarrow{g_1 \wedge g_2, a} E \{ \langle \langle \ell'_1, \ell'_2 \rangle, u_1 \cup u_2 \rangle \mapsto \nu_1(\langle \ell'_1, u_1 \rangle) \cdot \nu_2(\langle \ell'_2, u_2 \rangle) \mid \dots \}} \quad (\text{syn}),$$

plus a rule ind_2 for M_2 that is symmetric to ind_1 , with \dots in syn replaced by $\langle \ell'_1, u_1 \rangle \in \text{spt}(\nu_1) \wedge \langle \ell'_2, u_2 \rangle \in \text{spt}(\nu_2) \wedge u_1 \cup u_2$ is a function.

Inference rules ind_1 and ind_2 allow the individual VMDP to proceed independently if allowed by sync ; rule syn covers the case where they synchronise on a pair of actions. An element of sync is called a *synchronisation vector*; we also write $\langle a_1, a_2 \rangle \mapsto a$ for vector $\langle a_1, a_2, a \rangle$. This flexible form of parallel composition can be generalised to more than two VMDP with longer synchronisation vectors. We refer to such a general parallel composition as a *network* of VMDP and write e.g. $\|_{SV}(M_1, M_2, M_3)$ for the network of M_1 , M_2 , and M_3 , with the set of synchronisation vectors SV . The case in syn where $\nu_1(\langle \ell'_1, u_1 \rangle) \cdot \nu_2(\langle \ell'_2, u_2 \rangle) \neq 0$ but u_1 and u_2 assign different expressions to a shared variable (s.t. $u_1 \cup u_2$ is not a function) is a modelling error, and we do not consider networks of VMDP with such inconsistent assignments, or with inconsistent initial valuations x_{I_1} and x_{I_2} .

Example 3. Figure 3 shows an initial fragment of the network of the node VMDP N_1 through N_4 on the left. We omit location labels. Synchronisation vectors are

$$SV = \{ \langle \text{rcv}, \text{rcv}, \text{rcv}, \text{rcv} \rangle \mapsto \text{rcv}, \langle \text{snd}_1, \perp, \perp, \perp \rangle \mapsto \text{snd}_1, \dots \}:$$

In every slot, the nodes first independently make their choices of whether to send or not (snd_i vs. nop_i); then they synchronise on rcv to perform their

receive actions simultaneously. Updates are atomic, i.e. all right-hand sides are evaluated first, thus there is no order dependency in e.g. $\{c_3 \leftarrow c_3 + d, d \leftarrow 0\}$. On the right of Fig. 3, we show a fragment of the network’s MDP semantics. We write $\ell.v_d.v_{c_1}v_{c_2}v_{c_3}v_{c_4}$ for state $\langle \langle \dots, \ell, \dots \rangle, \{d \mapsto v_d, c_1 \mapsto v_{c_1}, \dots\} \rangle$.

A *path* defines the behaviour of an MDP by resolving all nondeterminism and probabilistic choices. A *scheduler* resolves the nondeterminism only.

Definition 5. *Let M be an MDP as in Definition 1. A path π of M is an infinite sequence $\pi = s_0 tr_0 s_1 \dots \in (S \times Tr(M))^\omega$ such that, for all $i \in \{0, \dots\}$, we have $tr_i = \langle a, \mu \rangle \in T(s_i)$ and $\mu(s_{i+1}) > 0$. $\Pi(M)$ is the set of all paths of M . We write $\Pi_{fin}(M)$ for the set of all path prefixes π_{fin} ending in a state. The last state of π_{fin} is denoted $last(\pi_{fin})$.*

Definition 6. *Let M be an MDP as above. A (randomised, history-dependent) scheduler is a function $\sigma: \Pi_{fin}(M) \rightarrow Dist(Tr(M))$ such that $\forall s \in S: \sigma(s) = tr \Rightarrow tr \in T(s)$. We write $\mathfrak{S}(M)$ for the set of all schedulers of M . A deterministic scheduler is in $\Pi_{fin}(M) \rightarrow Tr(M)$; a (deterministic) memoryless scheduler is in $S \rightarrow Tr(M)$. A memoryless scheduler σ_{ml} defines a corresponding deterministic scheduler σ_{det} by $\sigma_{det}(\pi_{fin}) = \sigma_{ml}(last(\pi_{fin}))$ and a deterministic scheduler σ_{det} defines a scheduler σ by $\sigma(\pi_{fin}) = \{\sigma_{det}(\pi_{fin}) \mapsto 1\}$.*

If we “apply” a scheduler σ to an MDP M , it removes all nondeterminism and we obtain a DTMC $M|_\sigma$, whose paths can be measured and assigned probabilities according to the transitions’ distributions. Formally, these probability measures over sets of paths are built via cylinder sets; we refer the interested reader to e.g. [25] for details. Given a set of goal states $G \subseteq S$ and a scheduler σ , we are interested in the probability of the measurable subset of paths of $\Pi(M|_\sigma)$ including a state in G . We call the supremum (infimum) when ranging over all schedulers $\sigma \in \mathfrak{S}(M)$ the *maximum (minimum) reachability probability*. There is always a memoryless scheduler that achieves the supremum respectively infimum [8].

Example 4. For our example contact plan, we are interested in the maximum probability to reach a state where N_4 has at least one copy, and the corresponding scheduler. Expanding the MDP of Example 3, we can calculate that probability to be 0.493. The optimal choices are to send in slots T_1 and T_2 ; then in slot T_3 ,

- if we are in state $3_a.0.1000$ (i.e. if the first copy was lost on the way from N_1 via N_2 to N_3): send from N_1 to N_3 , and send from N_3 to N_4 in T_4 ;
- if we are in state $3_a.0.1010$ (i.e. the first copy made it to N_3), do not send from N_1 to N_3 (since N_3 already has a copy and can only send to N_4 once), then send in both T_4 and T_5 .

Thus the optimal choice of node N_1 in slot T_3 depends on whether node N_3 has a copy of the message or not. In a real distributed setting, N_1 cannot know this.

3 Distributed Scheduling

Example 4 showed that a scheduler that maximises the probability for eventual message delivery may need complete *global* information. In a distributed system like a satellite constellation, such a scheduler cannot be implemented; the satellites must decide whether to send based on their local state (here: their number of message copies c_i) only. This problem was initially studied with a focus on aspects of compositionality [21], ignoring algorithms except for a simple partial-information setting [2]. Giro et al. [35] defined *distributed schedulers*, for which computing and approximating optimal probabilities is in general impossible [35]. The formalisms of [2] and [35] do not provide for scheduling the interleaving of parallel components, i.e. deciding which component acts first in case both of them have an enabled edge in the same state. This gap was filled in [36] along with the introduction of *strongly distributed schedulers*. Though model checking remains undecidable in general, [36] proved that for memoryless schedulers it is “only” NP-hard. Consequently, no model checker as of today supports (strongly) distributed schedulers. The only prototype [18] was restricted to time-bounded reachability and suffered from exponential explosion in intermediate model sizes. Other prototype tools only provide overapproximations [37,38]. Our formalism of Sect. 2 is more expressive than those previously considered by allowing interleaving and information exchange via both synchronisation and (shared) variables.

In this paper, we restrict to a memoryless deterministic variant of distributed schedulers adapted to our formalism (Sect. 3.1), and we define a desirable characteristic of models that makes scheduling decisions about interleavings irrelevant (Sect. 3.2). We prefer memoryless distributed schedulers because history-dependent ones need infinite memory (which is again unimplementable), and using randomised schedulers would add additional unpredictability to the behaviour of the system, which is undesirable from the practitioner’s point of view.

3.1 Simple Distributed Schedulers

Definition 7. *Given a network of VM DP $M = \parallel_{SV}(M_1, \dots, M_n)$ and $i \in \{1, \dots, n\}$, let $\text{read}(M_i)$ be the set of all variables that occur in the guards of edges of M_i or on the right-hand sides of assignments in the updates in M_i . A state in $\llbracket M \rrbracket = \langle S, s_I, A, T \rangle$ has the form $s = \langle \langle \ell_i, \dots, \ell_n \rangle, v \rangle$ where ℓ_i is the current location of M_i and $v \in \text{Val}_{\cup_i X_i}$. Then the M_i -projection of s is $s \downarrow_{M_i} \stackrel{\text{def}}{=} \langle \ell_i, v \downarrow_{\text{read}(M_i)} \rangle$ with $v \downarrow_{\text{read}(M_i)} \stackrel{\text{def}}{=} \{x \mapsto v(x) \mid x \in \text{read}(M_i)\} \in \text{Val}_{\text{read}(M_i)}$. Let $S \downarrow_{M_i} \stackrel{\text{def}}{=} \{s \downarrow_{M_i} \mid s \in S\}$ be the set of all projected states. Every transition $tr = s \xrightarrow{a} s'$ in $\llbracket M \rrbracket$ can be traced back to a unique (generalised) synchronisation vector $sv \in SV$ through the rules of Definitions 3 and 4. We say that M_i is involved in tr if $sv[i] \neq \perp$. We write $I_t(M_i)$ for the set of all transitions M_i is involved in. For a transition tr , $I_c(tr) \stackrel{\text{def}}{=} \{M_i \mid tr \in I_t(M_i)\}$ is the set of all components involved in tr , and for a set of transitions TR , $I_c(TR) \stackrel{\text{def}}{=} \bigcup_{tr \in TR} I_c(tr)$.*

Simple distributed schedulers now consist of an interleaving scheduler (to select the component to perform the next transition) plus a local scheduler per component that only sees the component's projection of the current state:

Definition 8. A simple distributed scheduler for M as above is a tuple $\sigma_{sd} = \langle \sigma_I, \sigma_1, \dots, \sigma_n \rangle$ of an interleaving scheduler $\sigma_I: S \rightarrow \mathbb{N}$ and n local schedulers $\sigma_i: S \rightarrow I_t(M_i) \cup \{\perp\}$ for $i \in \{1, \dots, n\}$ s.t. $\sigma_I(s) \in \{i \mid T(s) \cap I_t(M_i) \neq \emptyset\}$, $\sigma_{\sigma_I(s)}(s) \in T(s) \cap I_t(M_i)$, and $s \downarrow_{M_i} = s' \downarrow_{M_i} \Rightarrow \sigma_i(s) = \sigma_i(s')$ for all $s, s' \in S$, $i \in \{1, \dots, n\}$. It defines a memoryless scheduler σ for $\llbracket M \rrbracket$ by $\sigma(s) = \sigma_{\sigma_I(s)}(s)$.

Simple distributed schedulers differ from the partial-information setting of [2], and from hidden Markov models or partially-observable MDP, by combining multiple projections in one scheduler. They also differ from the (strongly) distributed schedulers of [35, 36] by hiding information in states but not admitting information disclosure via synchronisation (by virtue of being memoryless). As such, they match the nowadays standard state-based approach to probabilistic verification where transition labels are only used for synchronisation, as embedded in e.g. the JANI [14], MODEST [9, 40], and PRISM languages [46], and implemented in e.g. EPMC [41], the MODEST TOOLSET [43], PRISM [46], and STORM [24].

Unlike [18, 36], our formalism does not partition actions into inputs and outputs; thus every component involved in one of our (undirected) transitions can be chosen by the interleaving scheduler. Interleaving schedulers are problematic: They may disclose global information by scheduling transitions in certain orders (see [36]), the nondeterminism they deal with is in fact uncontrollable, and they would again be unimplementable in fully distributed systems.

3.2 Good-for-Distribution Models

In [36], the problem of information disclosure by the interleaving scheduler was solved by restricting to *strongly* distributed schedulers which do not reveal information in this way *by definition*. Since this is not a constructive approach, and since we cannot implement an interleaving scheduler anyway, we only create models where the interleaving scheduler is by construction irrelevant.

Definition 9. A network of VMDP $M = \parallel_{SV}(M_1, \dots, M_n)$ is good for distributed scheduling w.r.t. reachability of goal set G if in all states $s \in S$ of $\llbracket M \rrbracket = \langle S, s_I, A, T \rangle$ where $|T(s)| > 1 \wedge |\{i \mid T(s) \cap I(M_i) \neq \emptyset\}| > 1$ we have

$$\forall s \xrightarrow{a} s': s \in G \Leftrightarrow s' \in G \quad \vee \quad \forall s \xrightarrow{a} s': s \in G \Leftrightarrow s' \notin G, \quad (1)$$

$$\forall i \in \{1, \dots, n\}: |I_t(M_i) \cap T(s)| > 1 \Rightarrow I_c(I_t(M_i) \cap T(s)) = \{M_i\}, \quad (2)$$

$$\text{and} \quad s \xrightarrow{a} s' \Rightarrow \forall M_c \in \{M_1, \dots, M_n\} \setminus I_c(s \xrightarrow{a} s'): s' \downarrow_{M_c} = s' \downarrow_{M_c}. \quad (3)$$

In words, a network is good if in all states where the interleaving scheduler has a nontrivial choice among multiple components, (1) it cannot influence whether we directly move to a goal state, (2) no component has a local choice involving at least one synchronising transition, and (3) no transition can change variables that are visible to a component not involved in the transition.

Lemma 1. *For a good-for-distributed-scheduling network of VMDP M , fair interleaving schedulers σ_{I_1} and σ_{I_2} , and a set G of goal states, the maximum (minimum) probability to reach a state in G under σ_{I_1} is the same as under σ_{I_2} .*

An interleaving scheduler is *fair* if, on every cycle in $\llbracket M \rrbracket$, it chooses every available component at least once. This is a reasonable assumption in practice.

Proof Sketch. The conditions of Definition 9 only apply to states where the interleaving scheduler actually has a choice; in all others, σ_{I_1} and σ_{I_2} must coincide. So let s be a state where $\sigma_{I_1}(s) \neq \sigma_{I_2}(s)$ and let $TR_i = T(s) \cap I_i(M_{\sigma_{I_i}(s)})$ with $tr_i = s \xrightarrow{a_i} s'_i \in TR_i$ for $i \in \{1, 2\}$. Then $T(s'_1) \supseteq T(s) \setminus TR_i$ and $T(s'_2) \supseteq T(s) \setminus TR_i$, i.e. the choice of the interleaving scheduler cannot disable a non-local transition. This is because (a) condition 3 prevents the tr_i from changing the values of guards in non-involved components, and (b) condition 2 requires that, if more than one component is involved in tr_i , then $TR_i = \{tr_i\}$. The latter ensures that taking a synchronising transition cannot disable another synchronising transition by taking away a needed “synchronisation partner” (a transition with the same label as needed by the synchronisation vectors). Thus the restriction of $\llbracket M \rrbracket$ induced by an interleaving scheduler together with condition 1 and the fairness requirement is a (stronger) variant of a partial order reduction using ample sets as in [37] (with condition **A5'**).

The fairness requirement is trivially satisfied for acyclic models (modulo self-loops in leaf states) like our running example, or if we replace σ_I by σ_{uni} , the randomised scheduler that picks a component uniformly at random every time.

Example 5. Our running example from Examples 2 and 3 is good for distributed scheduling: in every state, either a single node internally decides between \mathbf{snd}_i and \mathbf{nop}_i or all nodes synchronise on \mathbf{rcv} . In the former states, the interleaving scheduler has no choice, thus no conditions apply. In the latter states, conditions 2 and 3 are directly satisfied. The only way to move into a goal state is from a state of the latter kind, thus condition 1 also holds.

While simple distributed schedulers restrict local schedulers from *reading* certain variables, a good-for-distributed-scheduling model restricts component edges from *writing* to certain variables. A model checker can easily determine whether a model is good, but a precise syntactic check on the network-of-VMDP level is not possible. We can still use syntactic overapproximations, e.g. by requiring that locations with multiple outgoing edges have only internal edges not writing to shared variables, and that a shared variable may only be updated by synchronising edges involving all components that read the variable. Our running example does not satisfy this syntactic restriction due to the writes to d on \mathbf{snd}_i -labelled edges. We solve this (in Sect. 5.1) by using a feature of MODEST and JANI that allows to specify *sequences* of atomic updates, moving the d -write assignments onto the \mathbf{rcv} transitions and executing them before the d -read assignments.

4 Lightweight Distributed Scheduler Sampling

Since PMC for distributed schedulers is undecidable or computationally infeasible, we propose a different approach to find useful high-probability strategies: we combine statistical model checking (SMC) [44, 56] with a new variant of lightweight scheduler sampling (LSS) [48] that samples only simple distributed schedulers. As our models are good for distributed scheduling by construction (see Sect. 5.1), only the satellites’ local choices are relevant and we can replace the interleaving scheduler by uniformly random choices.

SMC with LSS. SMC is Monte Carlo simulation with formal models: perform several *simulation runs* using a pseudo-random number generator (PRNG) to resolve probabilistic choices according to the model’s distributions, then return a statistical estimate and confidence for the probability of interest. As-is, SMC does not consider the optimisation problem over nondeterminism posed by MDP. LSS is to date the only extension of SMC that takes scheduling into account (in contrast to e.g. PRISM or UPPAAL SMC [23], which always use the uniformly random scheduler) and also preserves SMC’s constant memory usage (in contrast to learning-based approaches like [11]). The basic idea of LSS is as follows:

1. Randomly select m 32-bit integers. Each of them is a *scheduler identifier* σ .
2. For each σ , perform standard SMC under the scheduler identified by σ .
3. Return the maximum (or minimum) result and the corresponding σ .

Due to the multiple tests, we need to adjust SMC’s statistical evaluation [22, 48]. Within step 2, when there is a choice between n transitions from state s , LSS concatenates the bit-vector representations of s and σ into $s.\sigma$, hashes the result into a 32-bit number $h = \mathcal{H}(s.\sigma)$, and picks the $h \bmod n$ -th transition. \mathcal{H} is deterministic so that σ defines a fixed memoryless scheduler. If \mathcal{H} is also uniform (w.r.t. all bits of $s.\sigma$), then LSS uniformly samples memoryless schedulers.

The result of SMC with LSS is an *underapproximation* (overapproximation) of the maximum (minimum) probability. It can thus e.g. disprove safety, or show schedulability—which is what we are interested in. Its efficiency—how large an m we need to get a good approximation—is determined by the probability of sampling a near-optimal scheduler, and thus strongly depends on the model at hand.

LSS for Distributed Schedulers. In contrast to PMC, LSS is easy to adapt to different classes of schedulers by changing the input to \mathcal{H} . In Algorithm 1, we show pseudocode for our adaptation to simple deterministic schedulers for MDP. We write $\mathcal{U}(\mu)$ for the pseudo-random selection by PRNG \mathcal{U} of a value from $\text{spt}(\mu)$ according to the probabilities of μ . Line 5 implements the interleaving scheduler; we assume good-for-distributed-scheduling models and thus use uniform random resolution here, but could equally replace the line by $i := \mathcal{H}(\sigma.s) \bmod |T_i|$. Line 7 implements the local scheduler, whose input is restricted to the chosen component’s projection of the current state. In line 8, we use \mathcal{U}_{pr} to pseudo-randomly

Input: Network of VMDP $M = \parallel_{SV}(M_1, \dots, M_n)$ with $\llbracket M \rrbracket = \langle S, s_I, A, T \rangle$, goal set $G \subseteq S$, $\sigma \in \mathbb{Z}_{32}$, \mathcal{H} uniform deterministic, PRNG \mathcal{U}_{pr} .

```

1  $s := s_I$ 
2 while  $s \notin G$  do // break on goal state
3   if  $\forall s \xrightarrow{a} \mu: \mu = \{s \mapsto 1\}$  then break // break on self-loops
4    $C := \{j \mid T(s) \cap I_t(M_j) \neq \emptyset\}$  // get active components
5    $i := \mathcal{U}_{pr}(\{j \mapsto \frac{1}{|C|} \mid j \in C\})$  // select component uniformly
6    $T_i := T(s) \cap I_t(M_i)$  // get component's transitions
7    $\langle a, \mu \rangle := (\mathcal{H}(\sigma.s \downarrow_{M_i}) \bmod |T_i|)$ -th element of  $T_i$  // schedule local transition
8    $s := \mathcal{U}_{pr}(\mu)$  // select next state according to  $\mu$ 
9 return  $s \in G$ 

```

Algorithm 1. Lightweight simple distributed scheduler sampling

select the successor state according to the distribution determined by the scheduled transition. Line 3 terminates the simulation negatively if we find a state that only has deterministic self-loop transitions; this suffices for our space DTN models, but could be replaced by smarter loop detection or methods like [47, 55].

We have implemented Algorithm 1 in MODES [12], the statistical model checker of the MODEST TOOLSET [43]. MODES is implemented in C#, freely available at modestchecker.net, runs on 64-bit Linux, macOS, and Windows, and is faster than other current general-purpose SMC tools [13, Section 7.1]. Its input languages are MODEST [9, 40] and the tool-independent JANI model exchange format [14]. It provides both variants of line 5 discussed above, and implements corrected statistical tests as well as two-phase and smart sampling.




Fig. 4. Satellite DTN routing scheduling toolchain

5 Scheduling Satellite Communication

To apply our new LSS method of Sect. 4 to space DTN, we created the toolchain shown in Fig. 4. We use the STK tool by AGI [1] and the Contact Plan Designer plugin [27] to model the scenario and export the contact plan to a file in Interplanetary Overlay Network format (ION) [15]. This plan contains the precise real-time communication windows; we developed the Python CP2MODEST tool that, given such a plan, message source and destinations, and a bound on the number of copies, (1) abstracts the plan into the form of Fig. 1 with discrete

non-overlapping slots,¹ and (2) creates a MODEST model representing a network of VMDP of the same structure as Example 2. We then run MODES with LSS for simple distributed schedulers to obtain a good scheduler and its probability. Compared to the previous PMC-based work [51], we not only generate guaranteed implementable schedules, but also support multiple message copies.

5.1 Modelling Satellite DTN

The model of Fig. 1 given in Example 2 uses unidirectional unreliable communication: for every contact, one node is predetermined as sender; if communication fails, the copy is lost. We also assumed that there is at most one contact and we transmit at most one copy per slot. The models generated by CP2MODEST keep the same structure, but assume bidirectional half-duplex communication, support multiple contacts per slot (including one node having a contact with multiple others; it then needs to choose with whom to communicate), allow sending multiple copies (to one node in one slot), and use either unreliable or acknowledgment-based communication. The latter allows the sender to determine whether a message was successfully received, thus in such models it will keep the sent copies in case of failure. Unreliable communication is a natural choice in deep-space networks, the original application of DTN, while acknowledgment mechanisms are possible and typical in LEO constellations.

```

1  process Node1(int(0..COPIES) copies) {
2      alt {
3          :: nop1; rcv
4          :: when(copies >= 1) snd1to2_1;
5          rcv palt {
6              :0.9: {= data1 = 1, dest1 = 2, 2: copies -- ack2==1 ? 1 : 0 =}
7              :0.1: {= /* lost */ =} }
8          :: when(copies >= 2) snd1to2_2;
9          rcv palt {
10             :0.9: {= data1 = 2, dest1 = 2, 2: copies -- ack2==1 ? 2 : 0 =}
11             :0.1: {= /* lost */ =} }
12         :: rcv2to1;
13         rcv {= 1: copies += dest2==1 ? data2 : 0, 1: ack1 = 2 =}
14     };
15     rcv;
16     ... }

```

// slot 1: contact with node 2
// do nothing in this slot
// send one copy to node 2
// send two copies to node 2
// listen for communication from node 2
// slot 2: no contact
// three more slots of the same pattern

Listing 1. Excerpt of the MODEST code for node N_1 with a bound of 2 copies

The models created by CP2MODEST consist of one *process* definition per node. Listing 1 shows an excerpt of node N_1 in the plan of Fig. 1, but with acknowledgments and all mechanisms to support multiple contacts per slot in place. Like in the VMDP of Example 2, in every slot in which a node has a contact, we have a choice of action followed by a global synchronisation on action `rcv`. Here, in T_1 , the choices are to (a) do nothing (line 3), (b) try to send one (line 4) or two

¹ The abstraction underapproximates contacts (it may only remove or shorten communication opportunities), so a strategy for the abstract plan is always implementable.

copies (line 8) to N_2 , or (c) listen to N_2 (line 12). We use distinct actions for every choice to make it easier to trace the best scheduler’s decisions later on. Global variable `data1` takes the role of d of Example 2, but only for messages sent by N_1 , and `dest1` indicates the intended receiver—both of this is to support multiple contacts per slot. As mentioned at the end of Sect. 3.2, we moved these assignments onto the `rcv` edge to make the model syntactically good for distributed scheduling; assignments prefixed 2: are executed after prefix 1: which follow those with no prefix. The acknowledgment mechanism uses global variables like `ack2` indicating from whom node N_2 successfully received a message. The sending node then reduces its `copies` count by that number. In slots without a contact, like slot T_2 for N_1 , we just move to the next slot by a `rcv` together with the other nodes (line 15). Note that nodes cannot *create* copies; this is because they cannot know how many copies there are at other nodes and must not violate the global bound; we can let the initial node create as many copies as allowed w.l.o.g. since any number of copies can be transmitted in a slot.

5.2 The Walker Constellation

As a realistic case study, we propose a LEO satellite constellation in a Walker formation of three orbital planes each with four equally separated satellites. The constellation is thought to provide high-latency data service to ten isolated ground nodes randomly distributed around the globe. A ground station located in Córdoba, Argentina, provides an Internet gateway that the nodes access using DTN protocols. The ranges from ground station to satellites are set to 2000 km at a minimal elevation angle of 20° (mimicking a constrained antenna at the satellite). In turn, satellites can reach ground terminals at a distance of 1500 km at the same elevation angle. Figure 5 shows the ground track and a 3D visualisation of the constellation. Ground nodes (user terminals) are shown in red, the ground station in green, and satellites in cyan. The domes and cones over the ground nodes and under the satellites illustrate the communication ranges. The contact plan runs for 24 h from 01 Jul 2020 00:00:00.

5.3 Experiments

We applied our toolchain to the example contact plan of Fig. 1 and the Walker constellation described above, both with at most 2 message copies in the network. For the example contact plan with unreliable communication, we can easily determine the best simple distributed scheduler and its transmission probability by again expanding the MDP semantics of Fig. 3; the probability is 0.4645. Recall from Example 4 that global-information schedulers achieve probability 0.493. We can thus compare the probabilities computed by LSS with the actual values for distributed schedulers on this example. The state space of our model of the Walker constellation is still small enough for PMC, so we can compare the values obtained via PMC and LSS for global-information schedulers at least, but not for distributed schedulers due to the infeasibility of PMC in this case. Since the effectiveness of LSS depends on the rarity of near-optimal schedulers, we expect




Fig. 5. Visualisation of the Walker constellation

to see a tradeoff between LSS for global-information and LSS for distributed schedulers: There are many more global-information schedulers than distributed ones. So even though the former may realise higher probabilities, LSS might only rarely find any good global-information scheduler, whereas it may often find a good distributed scheduler due to their more limited choices.

Our experiments ran on an Intel Core i7-4790 workstation (3.6–4.0 GHz, 4 cores) with 64-bit Ubuntu Linux 18.04. We used smart sampling [22] with a fixed number of initial schedulers m_0 equal to the per-iteration budget n_0 . In iteration i , smart sampling performs $\lceil \frac{n_i}{m_i} \rceil$ simulation runs for each of the m_i schedulers, discards the “worst” half of the schedulers according to their current probability estimate, and moves to iteration $i + 1$ with $m_{i+1} = \lfloor \frac{m_i}{2} \rfloor$. We can thus cover a large number of schedulers with only $\approx \log_2(m_0) \cdot n_0$ simulation runs in total. We show the results in Table 1. We used the MODEST TOOLSET’s MCSTA model checker for PMC and MODES as described in Sect. 4 for the SMC-LSS- m_0 runs. Due to the randomised nature of the experiments, we repeated

each three times and report the average and highest maximum probability over the repetitions. Runtimes were at most 5 min, for the $m_0 = 100000$ setting, with distributed scheduler runs taking around 10% longer than those for global-information schedulers. We use the adaptive method of [20, Section III] for the statistical evaluation and request the probability for absolute error <0.0025 on the example and <0.005 on the Walker case to be 95%.

Table 1. Experimental results: average (highest) max. probabilities over 3 repetitions

model	PMC	SMC-LSS-1000		SMC-LSS-10000		SMC-LSS-100000	
	global	global	distrib.	global	distrib.	global	distrib.
example/unrel.	0.493	0.48 (0.49)	0.46 (0.47)	0.49 (0.49)	0.46 (0.47)	0.49 (0.49)	0.46 (0.46)
example/acks	0.505	0.49 (0.50)	0.46 (0.48)	0.50 (0.50)	0.50 (0.50)	0.50 (0.50)	0.50 (0.51)
walker/unrel.	0.438	0.03 (0.06)	0.21 (0.30)	0.10 (0.16)	0.30 (0.37)	0.26 (0.33)	0.37 (0.38)
walker/acks	0.734	0.36 (0.38)	0.47 (0.48)	0.38 (0.40)	0.54 (0.60)	0.45 (0.47)	0.54 (0.56)

We see that, on the example, both variants of LSS find respective optimal schedulers easily. With acknowledgment-based communication (“acks”), the optimal distributed scheduler curiously realises the same probability as the global one. We asked MODES to print traces under this scheduler identifier, and found that nodes N_1 and N_3 can collaborate to implement a strategy similar to the one of Example 4: In slot T_3 , N_1 *always* tries to send to N_3 . However, if N_3 already has a copy, it chooses `nop3` in this slot; then N_1 does not receive an acknowledgment and gets to keep its copy for sending in T_5 . While “sneaky”, such behaviour could clearly be implemented in satellites, and it being found validates the correctness and applicability of our distributed LSS approach. On the Walker case study, the tradeoff described earlier comes into play: distributed-scheduler LSS consistently finds better schedulers, and more reliably finds them even for lower values of m_0 , despite global-information schedulers being able to realise higher probabilities in principle. In particular, our new version of LSS consistently finds schedulers not too far from the maximum achievable with global information according to PMC—but its schedulers are guaranteed implementable.

6 Conclusion

We have developed new theory and tools to tackle the challenge of computing good and implementable routing strategies for satellite DTNs under uncertain contact plans, using formal methods technology. We have proposed a new, modern notion of distributed schedulers appropriate for the application, extended the LSS technique to work with this new notion, created a toolchain incorporating the new theory and technology, and applied it to a realistic case study. While LSS may be limited (e.g. by not being able to prove safety) in a general setting,

we showed that it works very well for this application—which particularly benefits from the flexibility of LSS w.r.t. handling different scheduler classes. Once we found a good scheduler, it can be implemented in a satellite by a program that feeds its identifier, the current slot offset, and the number of local copies into a reimplementaion of MODES’ hash function. We plan to adapt other methods (e.g. [10]) to extract compact human-readable descriptions of the scheduler, too.

Data Availability. All data generated in our experimental evaluation plus instructions to replicate the experiments are available at DOI [10.4121/uuid:6aa24e1a-3551-4073-b533-4ba6e408212d](https://doi.org/10.4121/uuid:6aa24e1a-3551-4073-b533-4ba6e408212d) [42].

References

1. AGI Systems Tool Kit (STK). <http://www.agi.com/STK>
2. de Alfaro, L.: The verification of probabilistic systems under memory less partial information policies is hard. In: 2nd International Workshop on Probabilistic Methods in Verification, pp. 19–32. Technical report CSR-99-8, University of Birmingham (1999)
3. de Alfaro, L., Henzinger, T.A., Jhala, R.: Compositional methods for probabilistic systems. In: Larsen, K.G., Nielsen, M. (eds.) CONCUR 2001. LNCS, vol. 2154, pp. 351–365. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44685-0_24
4. Alvarez, J., Walls, B.: Constellations, clusters, and communication technology: expanding small satellite access to space. In: 2016 IEEE Aerospace Conference, pp. 1–11 (2016). <https://doi.org/10.1109/AERO.2016.7500896>
5. Araniti, G., et al.: Contact graph routing in DTN space networks: overview, enhancements and performance. *IEEE Commun. Mag.* **53**(3), 38–46 (2015). <https://doi.org/10.1109/MCOM.2015.7060480>
6. Baier, C., de Alfaro, L., Forejt, V., Kwiatkowska, M.: Model checking probabilistic systems. *Handbook of Model Checking*, pp. 963–999. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_28
7. Baier, C., Katoen, J.P.: *Principles of Model Checking*. MIT Press, Cambridge (2008)
8. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Thiagarajan, P.S. (ed.) FSTTCS 1995. LNCS, vol. 1026, pp. 499–513. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-60692-0_70
9. Bohnenkamp, H.C., D’Argenio, P.R., Hermanns, H., Katoen, J.P.: MoDeST: a compositional modeling formalism for hard and softly timed systems. *IEEE Trans. Soft. Eng.* **32**(10), 812–830 (2006). <https://doi.org/10.1109/TSE.2006.104>
10. Brázdil, T., Chatterjee, K., Chmelík, M., Fellner, A., Křetínský, J.: Counterexample explanation by learning small strategies in Markov decision processes. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 158–177. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21690-4_10
11. Brázdil, T.: Verification of Markov decision processes using learning algorithms. In: Cassez, F., Raskin, J.-F. (eds.) ATVA 2014. LNCS, vol. 8837, pp. 98–114. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11936-6_8
12. Budde, C.E., D’Argenio, P.R., Hartmanns, A., Sedwards, S.: A statistical model checker for nondeterminism and rare events. In: Beyer, D., Huisman, M. (eds.) TACAS 2018. LNCS, vol. 10806, pp. 340–358. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89963-3_20

13. Budde, C.E., D'Argenio, P.R., Hartmanns, A., Sedwards, S.: An efficient statistical model checker for nondeterminism and rare events. STTT (2020, under review). <http://www.modestchecker.net/Publications/PDF/BDHS20-prelim.pdf>
14. Budde, C.E.: JANI: quantitative model and tool interaction. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 151–168. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_9
15. Burleigh, S.: Interplanetary overlay network: an implementation of the DTN bundle protocol. In: 4th IEEE Consumer Communications and Networking Conference, pp. 222–226 (2007). <https://doi.org/10.1109/CCNC.2007.51>
16. Burleigh, S.: Delay-tolerant networking: an approach to interplanetary Internet. IEEE Commun. Mag. **41**(6), 128–136 (2003). <https://doi.org/10.1109/MCOM.2003.1204759>
17. Caini, C., Firrincieli, R.: Application of contact graph routing to LEO satellite DTN communications. In: 2012 IEEE International Conference on Communications (ICC), pp. 3301–3305 (2012). <https://doi.org/10.1109/ICC.2012.6363686>
18. Calin, G., Crouzen, P., D'Argenio, P.R., Hahn, E.M., Zhang, L.: Time-bounded reachability in distributed input/output interactive probabilistic chains. In: van de Pol, J., Weber, M. (eds.) SPIN 2010. LNCS, vol. 6349, pp. 193–211. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16164-3_15
19. Cerf, V., et al.: Delay-tolerant networking architecture. RFC 4838, RFC Editor (April 2007). <http://www.rfc-editor.org/rfc/rfc4838.txt>
20. Chen, J., Xu, J.: Sampling adaptively using the Massart inequality for scalable learning. In: 12th International Conference on Machine Learning and Applications (ICMLA), pp. 362–367. IEEE (2013). <https://doi.org/10.1109/ICMLA.2013.149>
21. Cheung, L., Lynch, N., Segala, R., Vaandrager, F.: Switched probabilistic i/o automata. In: Liu, Z., Araki, K. (eds.) ICTAC 2004. LNCS, vol. 3407, pp. 494–510. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31862-0_35
22. D'Argenio, P., Legay, A., Sedwards, S., Traonouez, L.M.: Smart sampling for lightweight verification of Markov decision processes. STTT **17**(4), 469–484 (2015). <https://doi.org/10.1007/s10009-015-0383-0>
23. David, A., Larsen, K.G., Legay, A., Mikucionis, M., Poulsen, D.B.: Uppaal SMC tutorial. STTT **17**(4), 397–415 (2015). <https://doi.org/10.1007/s10009-014-0361-y>
24. Dehnert, C., Junges, S., Katoen, J.-P., Volk, M.: A Storm is coming: a modern probabilistic model checker. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 592–600. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63390-9_31
25. Forejt, V., Kwiatkowska, M., Norman, G., Parker, D.: Automated verification techniques for probabilistic systems. In: Bernardo, M., Issarny, V. (eds.) SFM 2011. LNCS, vol. 6659, pp. 53–113. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21455-4_3
26. Foust, J.: SpaceX's space-Internet woes: despite technical glitches, the company plans to launch the first of nearly 12,000 satellites in 2019. IEEE Spectr. **56**(1), 50–51 (2019). <https://doi.org/10.1109/MSPEC.2019.8594798>
27. Fraire, J.A.: Introducing contact plan designer: a planning tool for DTN-based space-terrestrial networks. In: 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT), pp. 124–127 (2017). <https://doi.org/10.1109/SMC-IT.2017.28>
28. Fraire, J.A., Burleigh, S., Finochietto, J.M.: Disruption-Tolerant Satellite Networks. ArtechHouse (2017)

29. Fraire, J.A., et al.: Assessing contact graph routing performance and reliability in distributed satellite constellations. *Hindawi J. Comput. Netw. Commun.* 2017, 18 p. (2017). Article ID 2830542. <https://doi.org/10.1155/2017/2830542>
30. Fraire, J.A., Madoery, P.G., Finochietto, J.M.: On the design and analysis of fair contact plans in predictable delay-tolerant networks. *IEEE Sens. J.* **14**(11), 3874–3882 (2014). <https://doi.org/10.1109/JSEN.2014.2348917>
31. Fraire, J.A., Nies, G., Gerstacker, C., Hermanns, H., Bay, K., Bisgaard, M.: Battery-aware contact plan design for LEO satellite constellations: the Ulloriaq case study. *IEEE Trans. Green Commun. Netw.* (2019). <https://doi.org/10.1109/TGCN.2019.2954166>
32. Fraire, J.A., Finochietto, J.M.: Design challenges in contact plans for disruption-tolerant satellite networks. *IEEE Commun. Mag.* **53**(5), 163–169 (2015). <https://doi.org/10.1109/MCOM.2015.7105656>
33. Fraire, J.A., Finochietto, J.M.: Routing-aware fair contact plan design for predictable delay tolerant networks. *Ad Hoc Netw.* **25**, 303–313 (2015). *New Research Challenges in Mobile, Opportunistic and Delay-Tolerant Networks Energy-Aware Data Centers: Architecture, Infrastructure, and Communication.* <https://doi.org/10.1016/j.adhoc.2014.07.006>
34. Fraire, J.A., Madoery, P.G., Finochietto, J.M., Leguizamón, G.: An evolutionary approach towards contact plan design for disruption-tolerant satellite networks. *Appl. Soft Comput.* **52**, 446–456 (2017). <https://doi.org/10.1016/j.asoc.2016.10.023>
35. Giro, S., D’Argenio, P.R.: Quantitative model checking revisited: neither decidable nor approximable. In: Raskin, J.-F., Thiagarajan, P.S. (eds.) *FORMATS 2007*. LNCS, vol. 4763, pp. 179–194. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75454-1_14
36. Giro, S., D’Argenio, P.R.: On the expressive power of schedulers in distributed probabilistic systems. *Electr. Notes Theor. Comput. Sci.* **253**(3), 45–71 (2009). <https://doi.org/10.1016/j.entcs.2009.10.005>
37. Giro, S., D’Argenio, P.R., Ferrer Fioriti, L.M.: Partial order reduction for probabilistic systems: a revision for distributed schedulers. In: Bravetti, M., Zavattaro, G. (eds.) *CONCUR 2009*. LNCS, vol. 5710, pp. 338–353. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04081-8_23
38. Giro, S., Rabe, M.N.: Verification of partial-information probabilistic systems using counterexample-guided refinements. In: Chakraborty, S., Mukund, M. (eds.) *ATVA 2012*. LNCS, pp. 333–348. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33386-6_26
39. Gottzein, E.: Challenges in the control and autonomy of communications satellites. *Control Eng. Pract.* **8**(4), 409–427 (2000). [https://doi.org/10.1016/S0967-0661\(99\)00171-9](https://doi.org/10.1016/S0967-0661(99)00171-9)
40. Hahn, E.M., Hartmanns, A., Hermanns, H., Katoen, J.P.: A compositional modelling and analysis framework for stochastic hybrid systems. *Formal Meth. Syst. Des.* **43**(2), 191–232 (2013). <https://doi.org/10.1007/s10703-012-0167-z>
41. Hahn, E.M., Li, Y., Schewe, S., Turrini, A., Zhang, L.: *ISCASMC*: a web-based probabilistic model checker. In: Jones, C., Pihlajasaari, P., Sun, J. (eds.) *FM 2014*. LNCS, vol. 8442, pp. 312–317. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06410-9_22
42. Hartmanns, A.: Sampling distributed schedulers for resilient space communication (artifact). 4TU.Centre for Research Data (2020). <https://doi.org/10.4121/uuid:6aa24e1a-3551-4073-b533-4ba6e408212d>

43. Hartmanns, A., Hermanns, H.: The Modest Toolset: an integrated environment for quantitative modelling and verification. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 593–598. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54862-8_51
44. Hérault, T., Lassaigne, R., Magniette, F., Peyronnet, S.: Approximate probabilistic model checking. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 73–84. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24622-0_8
45. Jenkins, A., Kuzminsky, S., Gifford, K.K., Pitts, R.L., Nichols, K.: DTN: flight test results from the international space station. In: 2010 IEEE Aerospace Conference, pp. 1–8 (2010)
46. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_47
47. Lassaigne, R., Peyronnet, S.: Probabilistic verification and approximation. *Ann. Pure Appl. Logic* **152**(1–3), 122–131 (2008). <https://doi.org/10.1016/j.apal.2007.11.006>
48. Legay, A., Sedwards, S., Traonouez, L.-M.: Scalable verification of Markov decision processes. In: Canal, C., Idani, A. (eds.) SEFM 2014. LNCS, vol. 8938, pp. 350–362. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15201-1_23
49. Madoery, P.G., Raverta, F.D., Fraire, J.A., Finochietto, J.M.: Routing in space delay tolerant networks under uncertain contact plans. In: 2018 IEEE International Conference on Communications (ICC), May 2018, pp. 1–6 (2018). <https://doi.org/10.1109/ICC.2018.8422917>
50. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. Wiley, Hoboken (1994)
51. Raverta, F.D., Demasi, R., Madoery, P.G., Fraire, J.A., Finochietto, J.M., D'Argenio, P.R.: A Markov decision process for routing in space DTNs with uncertain contact plans. In: 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE), pp. 189–194. IEEE (2018). <https://doi.org/10.1109/WiSEE.2018.8637330>
52. Sahai, A., Tandra, R., Mishra, S.M., Hoven, N.: Fundamental design tradeoffs in cognitive radio systems. In: Proceedings of the 1st International Workshop on Technology and Policy for Accessing Spectrum, p. 2. ACM (2006)
53. Sheng, M., Xu, G., Fang, X.: The routing of interplanetary Internet. *China Commun.* **3**(6), 63–73 (2006)
54. Wyatt, J., Burleigh, S., Jones, R., Torgerson, L., Wissler, S.: Disruption tolerant networking flight validation experiment on NASA's EPOXI mission. In: First International Conference on Advances in Satellite and Space Communications (SPACOMM), pp. 187–196 (2009). <https://doi.org/10.1109/SPACOMM.2009.39>
55. Younes, H.L.S., Clarke, E.M., Zuliani, P.: Statistical verification of probabilistic properties with unbounded until. In: Davies, J., Silva, L., Simao, A. (eds.) SBMF 2010. LNCS, vol. 6527, pp. 144–160. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19829-8_10
56. Younes, H.L.S., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 223–235. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45657-0_17