Centralized and Distributed Approaches for Restoring the Weak Controllability of Multi-Agent Interdependent STNUs

Extended Abstract

Ajdin Sumic
DTIS, ONERA, Université de Toulouse
Toulouse, France
aidin.sumic@onera.fr

Thierry Vidal
LAAS
Toulouse, France
thierry.vidal@laas.fr

Gauthier Picard
DTIS, ONERA, Université de Toulouse
Toulouse, France
gauthier.picard@onera.fr

Frédéric Maris
IRIT, Université de Toulouse
Toulouse, France
frederic.maris@irit.fr

Roberto Posenato
University of Verona
Verona, Italy
roberto.posenato@univr.it

Carlo Combi
University of Verona
Verona, Italy
carlo.combi@univr.it

ABSTRACT

This paper models the negative cycles that cause uncontrollability in Multi-Agent Interdependent Simple Temporal Networks with Uncertainty (MISTNU) as linear constraints, enabling (1) a fast centralized linear-programming repair and (2) a novel distributed constraint-reasoning approach that treats those cycles as inter-agent DCOP constraints, preserving privacy, with a performance comparison between distributed solvers and a centralized baseline.

KEYWORDS

Multi-agent planning, Temporal controllability, DCOP modeling

ACM Reference Format:
Ajdin Sumic, Thierry Vidal, Gauthier Picard, Frédéric Maris, Roberto Posenato, and Carlo Combi. 2026. Centralized and Distributed Approaches for Restoring the Weak Controllability of Multi-Agent Interdependent STNUs: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25–29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/FQWT7513

1 INTRODUCTION AND RELATED WORK

In multi-agent temporal planning, agents first build individual plans through task allocation and planning (possibly via negotiation) and then must check and coordinate their temporal executability. We focus on this second stage in a setting where the task owner controls task durations, while other compliant agents may depend on them and learn their values only near or during execution. Besides checking executability, such agents may *repair* a non-executable plan by negotiating shared constraints.

A Simple Temporal Network (STN) supports temporal consistency checking [2]. STNUs extend STNs with uncertain (*contingent*) durations and replace consistency with *controllability* [12]: a plan is controllable if an execution strategy exists regardless of contingent outcomes. Controllability has three levels—Weak, Dynamic, and Strong—depending on when contingents become known.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25–29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/FQWT7513

For distributed coordination, the Multi-Agent Interdependent STNUs model (MISTNU) [9] lets agents plan locally while synchronizing through commitments. It distinguishes exogenous contingencies from inter-agent contingencies (*contracts*), which are controlled by one agent but contingent for others. Prior work defined distributed controllability checking for MISTNUs and proposed a centralized SMT-based repair that renegotiates contract durations [1].

Weak and Dynamic uncontrollabilities in STNUs are characterized by negative cycles [5, 10]. We show that exploiting these cycles enables more targeted ("informed") repairs than the existing "blind" centralized approach. Our first contribution translates negative cycles into linear constraints that contain only contract information for the *Weak Repair* problem, and solves it via centralized linear programming. Our second contribution presents a distributed approach that satisfies the same constraints, modeling repair as a Distributed Constraint Optimization Problem (DCOP) [7, 13].

2 MULTI-AGENT INTERDEPENDENT STNUS

STNUs have been extended to multi-agent settings, where contingent links in one agent's network may depend on decisions made by other agents [11]. These shared, duplicated constraints are called *contracts*: they are controllable for the agent making the decision (the owner) but contingent for dependent agents (readers). To capture this, our previous work has introduced *Contracting* STNUs (cSTNUs) by associating each contract label p with a duration interval that can be attached to duplicated binary constraints across different STNUs [12]. Contract bounds and ownership are defined externally and shared among agents; during controllability checking, contracts are treated as contingent commitments that the owner is not expected to tighten, while other agents verify controllability with respect to them.

Then, a MISTNU is a model $\mathcal{G} = \langle A, \Sigma, B \rangle$ composed of a set of agents A, a set of cSTNUs Σ where each agent has its own cSTNU, and might own or observe some contracts; some of them being negotiable[1]. The cSTNUs are synchronized through a reference timepoint v_z (i.e., a specific time). Then, $\mathcal{G}$ is also composed of a reification function B that, given a contract p, returns the interval

[1]A contract is non-negotiable if it comes from *Nature* and hence has no owner. In this paper, we assume such contracts do not exist.

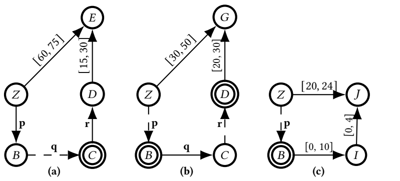

Contract	B	B'
p	[5, 10]	**[10, 10]**
q	[10, 15]	**[10, 10]**
r	[10, 20]	**[10, 10]**

Figure 1: The corresponding MISTNU of Example 1. Node Z acts as v_z. Doubly circled nodes are contingent timepoints.

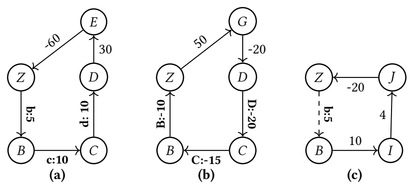

Figure 2: Inconsistent cycles of agent STNUs. STNU (a) is not WC as there is a cycle of length $\alpha = -5$.

$[l_p, u_p]$ of possible durations for p. Any cSTNU can be refined into an STNU by replacing contracts' labels with the interval $[l_p, u_p]$ delimited by B. A MISTNU is Weakly Controllable (WC) if each cSTNU is Weakly Controllable (which is enforced by checking the WC of their reification into an STNU). If not, it may be repaired by negotiating the durations of negotiable contracts, i.e., shrinking their bounds. We tackle the case of Weak Controllability, which can be relevant in applications where, for example, an agent's resource availability is not known in advance, but is determined along with related task durations at the start of the day [10].

Example 1. *Consider a radiologist (agent a), a nurse (agent b), and a doctor (agent c). The radiologist performs an X-ray on a patient (requirement $Z \xrightarrow{p} B$), after which the doctor analyzes the results to prescribe medication (requirement $B \xrightarrow{[0,10]} I$). During the X-ray, the nurse waits for it to finish (contingent $Z \dashrightarrow^{p} B$), escorts the patient to the waiting room (requirement $B \xrightarrow{q} C$), and brings a second patient to the radiologist (contingent $C \dashrightarrow^{r} D$). Hence, the scenario includes three contracts: p (first X-ray), q (escort/turnover), and r (second patient/X-ray coordination). After the second X-ray (requirement $C \xrightarrow{r} D$), the radiologist cleans and inspects the equipment (requirement $D \xrightarrow{[15,30]} E$). Figure 1 shows the corresponding MISTNU, which is not WC, with reification function B, and a repair is given with B′.*

3 CENTRALIZED WEAK REPAIR

In this section, we propose a (centralized) algorithm that improves the work in [11] to solve the Weak Repair Problem. First, for WC checking, we use the algorithm in [10], which returns all causes of non-controllability as inconsistent cycles.

An inconsistent cycle refers to a negative cycle in the distance graph of an STNU [5], which is a sequence of timepoints and edge values such that its length α is negative. The length of a cycle is the sum of the values of the edges composing the path. For example, each agent's network in the MISTNU depicted in Figure 1 has a negative cycle as illustrated in Figure 2.

Repairing a non-WC MISTNU is equivalent to shrinking the bounds of contracts such that all inconsistent cycles found in all the networks are repaired, i.e., their length becomes positive. We express an inconsistent cycle as a linear expression with only the bounds of the contracts involved in the inconsistent cycle, forcing the new bounds (l^p/u^p) to recover at least α time units to repair the inconsistent cycle. For example, the inconsistent cycle of agent a in Figure 2 is encoded as $(l^p - 5) + (l^q - 10) + (l^r - 10) \geq 5$, and a possible solution is $l^p = 10, l^q = 10$, and $l^r = 10$. Consequently, our new centralized repair algorithm looks for all inconsistent cycles and repairs the associated set of linear constraints[2]. An optimization function is added to find the minimal reduction in contracts.

4 DISTRIBUTED WEAK REPAIR SOLVING

In this section, WC checking is now performed in a distributed manner: each agent checks WC locally and, if needed, identifies inconsistent cycles. Agents then share cycles involving other agents' contracts, so each owner learns the constraints their contracts must satisfy. If the agents agree on contract bounds that satisfy all such constraints, the MISTNU is WC. Privacy is preserved because cycles are expressed with only the bounds of the contracts, which are public information.

We transformed the distributed Weak Repair Problem of MISTNU into a Distributed Constraint Optimization Problem (DCOP) [8]. The set of agents, the variables (l^p/u^p), and their owners are provided by the MISTNU. The domain of each variable is given by B, assuming a discretization $\{l_p, \ldots, u_p\}$ of the contracts' interval, without loss of generality. Finally, the set of weighted constraints can be constructed by transforming each linear constraint into a weighted constraint equal to 0 if the constraint is satisfied, $+\infty$ otherwise. To find optimal solutions, we add one weighted constraint per contract, where the weight is equal to the width of the new contract bounds: $u^p - l^p$. Now, because DCOP solvers will attempt to find a solution that minimizes the sum of the costs of the weighted constraints, they will either find one minimizing the reduction in the contracts or an invalid repair with a cost of $+\infty$. The latter means that no repair solution exists.

5 EXPERIMENTS AND DISCUSSION

We tested our methods on the benchmark of [11], seeking optimal solutions within a one-hour time limit. Our centralized method solved 2.5 more instances than the SMT algorithm presented in [11] (250 vs 100), despite an exponential number of inconsistent cycles to repair. We tested our distributed approach with different optimal DCOP solvers: SyncBB [4], DPOP [8], ADOPT [7], and AFB [3]. Only SyncBB and AFB performed well here, while the other two failed. Both performed only slightly better than the centralized SMT algorithm. Still, we observed particularly high-quality solutions when non-optimal solutions are considered (all instances solved within a few seconds), providing a practical trade-off.

Our method could be extended to Dynamic Controllability (DC) by repeatedly checking DC and repairing inconsistent cycles until DC is guaranteed or no solution exists that repairs all the inconsistent cycles found. That is due to the fact that current checking algorithms return only one inconsistent cycle at a time [6].

[2] Please note that additional constraints are added to ensure $l_p \leq l^p \leq u^p \leq u_p$.

REFERENCES

[1] Clark Barrett and Cesare Tinelli. 2018. *Satisfiability modulo theories.* Springer.

[2] Rina Dechter, Itay Meiri, and Judea Pearl. 1991. Temporal constraint networks. *Artificial intelligence* 49, 1-3 (1991), 61–95.

[3] Amir Gershman, Amnon Meisels, and Roie Zivan. 2006. Asynchronous Forward-Bounding for Distributed Constraints Optimization. In *Proceedings of the 2006 Conference on ECAI 2006: 17th European Conference on Artificial Intelligence August 29 – September 1, 2006, Riva Del Garda, Italy.* IOS Press, NLD, 103–107.

[4] Katsutoshi Hirayama and Makoto Yokoo. 1997. Distributed Partial Constraint Satisfaction Problem. In *Principles and Practice of Constraint Programming - CP97, Third International Conference, Linz, Austria, October 29 - November 1, 1997, Proceedings (Lecture Notes in Computer Science, Vol. 1330)*, Gert Smolka (Ed.). Springer, 222–236. https://doi.org/10.1007/BFB0017442

[5] Luke Hunsberger and Roberto Posenato. 2022. Speeding Up the RUL^{-} Dynamic-Controllability-Checking Algorithm for Simple Temporal Networks with Uncertainty. In *36th AAAI Conf. on Artificial Intelligence 2022.* 9776–9785. https://doi.org/10.1609/aaai.v36i9.21213

[6] Luke Hunsberger and Roberto Posenato. 2025. Recent algorithmic advances in simple temporal networks with uncertainty: From faster controllability checking to faster execution. *Inf. Comput.* 307 (2025), 105356. https://doi.org/10.1016/J.IC.2025.105356

[7] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. 2005. ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees. *Artificial Intelligence* 161, 1–2 (2005), 149–180.

[8] Adrian Petcu and Boi Faltings. 2005. A Scalable Method for Multiagent Constraint Optimization. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005*, Leslie Pack Kaelbling and Alessandro Saffiotti (Eds.). Professional Book Center, 266–271. http://ijcai.org/Proceedings/05/Papers/0445.pdf

[9] Ajdin Sumic, Alessandro Cimatti, Andrea Micheli, and Thierry Vidal. 2024. SMT-Based Repair of Disjunctive Temporal Networks with Uncertainty: Strong and Weak Controllability. In *21st International Conference, CPAIOR 2024 (Lecture Notes in Computer Science).* Springer.

[10] Aïdin Sumic and Thierry Vidal. 2024. A more efficient and informed algorithm to check Weak Controllability of Simple Temporal Networks with Uncertainty. In *18èmes Journées d'Intelligence Artificielle Fondamentale et 19èmes Journées Francophones sur la Planification, la Décision et l'Apprentissage pour la conduite de systèmes (JIAF-JFPDA'24).*

[11] Ajdin Sumic, Thierry Vidal, Andrea Micheli, and Alessandro Cimatti. 2024. Introducing Interdependent Simple Temporal Networks with Uncertainty for Multi-Agent Temporal Planning. In *31st International Symposium on Temporal Representation and Reasoning.*

[12] Thierry Vidal and Hélène Fargier. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. *Journal of Experimental & Theoretical Artificial Intelligence* 11, 1 (1999), 23–45.

[13] Makoto Yokoo, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara. 1998. The Distributed Constraint Satisfaction Problem: Formalization and Algorithms. *IEEE Transactions on Knowledge and Data Engineering* 10, 5 (1998), 673–685.