

Received: 20 April 2026 / Accepted: 26 May 2026 / Published online: 07 June 2026

*digital twin architecture,
immature processes, automated planning
simulation-based validation*

Alexander BOTT^{1*}, Suliot SALIHI¹,
Alexander PUCHTA¹, Jürgen FLEISCHER¹

PLANNING-BASED DECISION SPACE EXPLORATION FOR DIGITAL TWINS IN IMMATURE PRODUCTION PROCESSES

Highly variant and immature production processes frequently occur during the ramp-up of new manufacturing technologies, where process parameters, resource configurations, and operating strategies must be established under incomplete process knowledge. The resulting combinatorial configuration spaces make systematic planning and evaluation of alternative process setups difficult. This paper proposes a planning-based digital twin architecture that enables structured exploration and reduction of such decision spaces while integrating simulation-based feasibility assessment. The approach combines formal modeling of discrete process parameter variants with automated planning and physics-based simulation using PyBullet. Symbolic planning operates on an abstract representation of process steps, resources, and parameter variants to generate consistent process paths under cost and quality objectives. These process paths are subsequently instantiated as executable simulation scenarios, allowing verification of their physical feasibility and operational behavior. The architecture integrates decision generation, execution, and evaluation within a unified and modular pipeline. A virtual thermoforming demonstrator is used to verify the functional feasibility of the approach and to illustrate systematic decision space reduction through constraint-based planning. The results demonstrate that planning-based digital twin architectures provide a scalable foundation for supporting decision-making during ramp-up and configuration of highly variant production processes.

1. INTRODUCTION

Modern manufacturing increasingly faces the challenge of configuring production systems under uncertainty, especially during ramp-up phases of new technologies where product designs, materials, and process chains evolve rapidly [1]. A prominent example is electric mobility, where shorter product life cycles, an increasing number of variants, and technological uncertainties make production planning and system design substantially more challenging than in stable, long-running series production [2–4]. In such environments, decision-making is often experience-driven and reactive, which limits systematic process understanding, scalability, and robustness [5, 6]. In addition, production systems face the need

¹ wbk - Institute of Production Science, Karlsruhe Institute of Technology, Germany

* E-mail: alexander.bott@kit.edu

<https://doi.org/10.36897/jme/222460>

to adapt more frequently due to rising variant diversity and shorter life cycles, which increases the importance of structured scenario analysis and decision support [7, 8].

Digital twins provide a promising basis for supporting such processes by linking virtual representations with physical systems through continuous data exchange. In contrast to isolated digital models, they enable simulation-based analysis and virtual safeguarding of decisions before physical realization [9]. Enabled by advances in Industry 4.0 technologies, digital twins are increasingly integrated into operational manufacturing systems and support improvements in quality, cost efficiency, and system availability [10–13]. Within the conceptualization of digital twins by Eramo et al. [14], decision-support functionality can be realized through a combination of predictive and prescriptive models, where prescriptive elements drive the execution of selected strategies.

Despite these potentials, the key challenge in immature and highly variant environments is often not only simulating process behaviour, but structuring and exploring large combinatorial decision spaces under incomplete knowledge. This becomes particularly visible when many discrete product and process parameters interact and feasible configurations must be derived under changing boundary conditions. While classical optimization or manual parameter studies can be applied in principle, their systematic use becomes increasingly difficult as the number of coupled choices and constraints grows.

This paper proposes a planning-based digital twin architecture for structured decision space exploration in highly variant production processes. The approach combines formal modeling of discrete process parameter variants, automated planning, and PyBullet-based simulation to generate consistent, goal-oriented process configurations under cost and quality considerations and to verify their consistent transfer into executable simulation scenarios. In the terminology of Eramo et al. [14], the planning component can be interpreted as a prescriptive model that derives actionable configurations based on formal constraints and evaluation feedback. The architecture is demonstrated using a virtual thermoforming demonstrator, illustrating how constraint-based planning enables systematic reduction and exploration of large configuration spaces.

2. STATE OF THE ART

2.1. SIMULATION-BASED PROCESS EVALUATION

Simulation-based methods are widely used to analyse and evaluate production processes without interfering with real systems. In manufacturing, classical approaches predominantly rely on discrete-event simulation, in which process steps are represented as event-driven state transitions. Such methods are well suited for material flow analysis, bottleneck identification, and layout planning [15]. However, they provide only limited insight into physical interactions such as forces, contacts, or motion dynamics.

At the other end of the spectrum, high-fidelity numerical methods, most notably the finite element method (FEM), enable detailed analysis of physical effects but are computationally expensive and therefore unsuitable for large-scale exploration of alternative process configurations [15]. Physics-based simulation environments represent an inter-

mediate approach by approximating system behaviour through rigid-body dynamics, kinematic modelling, and contact formulations, thereby balancing physical realism and computational efficiency [16, 17].

PyBullet exemplifies such rigid-body physics simulation frameworks, providing access to rigid-body dynamics, collision detection, and contact modelling through a programmable interface [18]. Mechanical structures and robotic systems are typically defined using standardized formats such as URDF, enabling structured modelling of kinematic chains and interactions between resources and workpieces [19]. Recent extensions such as PyBullet Industrial further integrate robot simulation with simplified manufacturing process models, enabling the joint simulation of robotic manipulators and manufacturing operations within a single environment [20]. This allows the investigation of interactions between robots, tools, and workpieces, including processes that add, remove, or manipulate material. In production contexts, physics-based simulation therefore complements formal reasoning by enabling the evaluation of physically plausible execution behaviour prior to real-world deployment, thereby serving as an execution-level validation mechanism for alternative process configurations. To enable systematic exploration of such configurations, the decision space exploration problem is formalized as a planning problem defined by the tuple $(S, A, \gamma, s_0, G, J)$, where:

- S denotes the set of system states defined by parameter configurations,
- A the set of process actions,
- γ the state transition function,
- s_0 the initial state,
- G the goal condition,
- J the objective function.

2.2. FORMAL MODELING AND PLANNING APPROACHES

Formal modeling and symbolic planning form the decision-making layer of the proposed architecture. While physics-based simulation enables the evaluation of physically plausible execution behaviour, formal models are required to systematically structure the underlying decision space and enable automated reasoning over alternative configurations.

In manufacturing research, Product–Process–Resource models are commonly used to represent relationships between system entities. Building on this concept, the Product–Process–Resource–Skill (PPRS) model extends classical representations by explicitly introducing skills as the link between abstract processes and concrete resource capabilities [21]. Processes describe product-independent production activities, while skills capture the ability of specific resources to execute these processes. Tasks arise when skills are applied to concrete product instances.

While PPRS provides a structural description of variability, systematic exploration of large combinatorial decision spaces requires a formal reasoning mechanism. Symbolic planning addresses this need by operating on formally defined states, actions, and goal conditions [22]. The Planning Domain Definition Language (PDDL) has emerged as a standard formalism for representing such planning problems [23]. A PDDL domain specifies

the general action space and system dynamics, while a PDDL problem instantiates this structure with concrete objects, an initial state, and goal conditions [24]. As emphasized by Purohit and Chouhan [25], only the combination of domain and problem defines a concrete planning task. This separation supports reuse of formal models across scenarios and enables systematic exploration of alternative configurations within explicitly defined constraints.

2.3. DIGITAL TWINS IN PRODUCTION SYSTEMS

Digital twins are widely recognized as a key enabler for analysing and optimizing complex production systems. Conceptually, a digital twin represents a virtual counterpart of a physical system that is continuously synchronized via data exchange, allowing the virtual instance to reflect the operational state of its physical counterpart [10, 26]. This bidirectional coupling distinguishes digital twins from isolated simulation models.

In manufacturing environments, digital twins are embedded in cyber-physical production systems and Industry 4.0 architectures, integrating sensor infrastructures, IoT connectivity, and data-driven services [9, 27]. Typical applications include simulation-based process optimization, virtual commissioning, quality prediction, and predictive maintenance [12]. A commonly used differentiation distinguishes digital models, digital shadows, and digital twins based on data connectivity, with digital twins enabling closed feedback loops between virtual and physical domains [29].

Despite their widespread adoption, most industrial digital twin implementations focus on monitoring, lifecycle management, or optimization of comparatively stable production processes. In highly variant and immature manufacturing environments, however, the central challenge lies in structuring and exploring large combinatorial configuration spaces prior to physical realization. While digital twin architectures provide data integration and simulation capabilities, explicit mechanisms for formal decision modeling and systematic configuration exploration remain limited. This observation motivates architectures that tightly integrate formal modeling, planning, and simulation to support decision space exploration.

3. PLANNING-BASED DIGITAL TWIN ARCHITECTURE

Highly variant manufacturing processes often involve numerous interdependent decisions regarding process steps, parameter settings, and resource configurations. In the investigated thermoforming scenario, for example, multiple production steps such as gripping, clamping, transfer, heating, and forming can be executed using different parameter settings, tool configurations, and resource assignments. Even for a single product, this results in a large number of feasible process paths whose physical feasibility and performance characteristics are difficult to assess manually. The proposed architecture addresses this challenge by enabling a structured exploration of such high-dimensional decision spaces. Instead of explicitly enumerating all possible process configurations, it integrates formal decision space modeling, automated planning, physics-based simulation, and systematic evaluation in a unified pipeline. In contrast to digital twin approaches that

focus primarily on monitoring or isolated simulation, the presented system is designed as a planning-capable and feedback-driven architecture that combines decision generation, execution, and evaluation within a closed-loop framework. Figure 1 illustrates the proposed architecture and its modular structure. The architecture can be interpreted as a sequence of transformations:

Decision space $D \rightarrow$ planning:

$P(D)=\pi$

Plan $\pi \rightarrow$ simulation model:

$f(\pi)=M_{sim}$

Simulation execution \rightarrow observations:

$E(M_{sim})=O$

Evaluation \rightarrow objective values:

$g(O)=J(\pi)$

Importantly, feasibility is ensured intrinsically within the planning process. In the PDDL-based formulation, constraints are encoded directly through action preconditions and state transition rules. As a result, only valid process paths are generated by the planner, and the objective function $J(\pi)$ is minimized exclusively over this set of constraint-satisfying solutions. To evaluate the effectiveness of the proposed approach, the following hypotheses are formulated:

- H1: Planning-based decision space exploration achieves a significantly greater reduction of the feasible search space compared to random sampling.
- H2: Planning-based decision generation results in a higher proportion of physically feasible simulation outcomes than both random sampling and rule-based filtering.

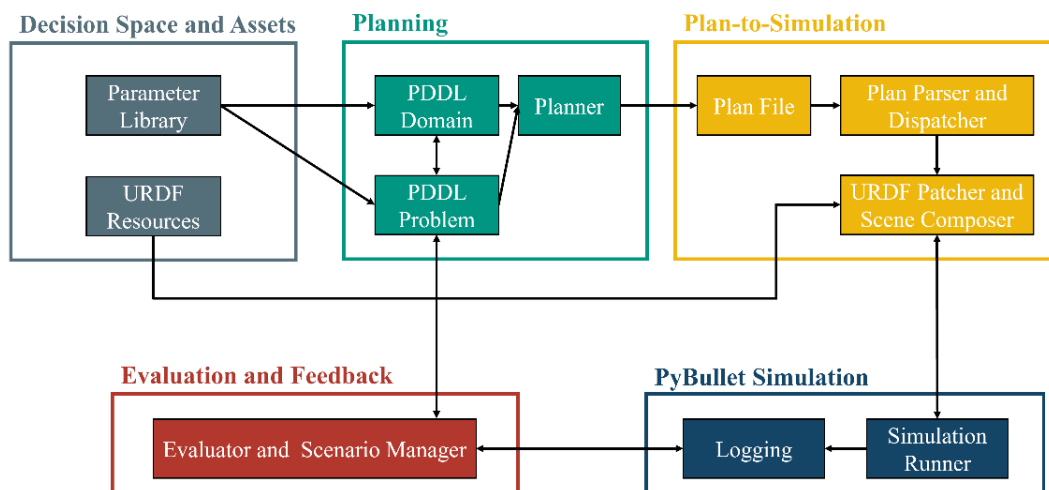


Fig. 1. Overview of the proposed planning-based digital twin architecture integrating decision space modelling, symbolic planning, simulation, and evaluation in a closed-loop framework.

The Decision Space and Assets form the foundation of the architecture. The decision space is derived from a structured definition of process steps, process parameters, and discrete process parameter variants, which are organized in a parameter library. This library provides a machine-interpretable representation of admissible decision alternatives and structural

constraints. Parameters are modeled discretely to ensure compatibility with symbolic planning and to retain explicit control over combinatorial complexity. Complementing this abstract representation, geometric and physical assets describe the executable entities of the system, including tools, fixtures, manipulators, and workpieces together with their structural and physical properties. While the parameter library governs logical decision-making, the asset models enable physically consistent execution during simulation. The strict separation between abstract decision parameters and executable assets supports extensibility and allows additional variants to be integrated without architectural modifications.

The Planning component constitutes the central transformation layer between the structured decision space and executable process behaviour. A process plan

$$\pi = (a_1, \dots, a_n) \quad (1)$$

is defined as a sequence of actions $a_i \in A$ that transforms an initial state s_0 into a goal state s_g , with state transitions given by

$$s_{i+1} = \gamma(s_i, a_i) \quad (2)$$

Planning is formulated symbolically, with process steps represented as actions characterized by explicit preconditions and effects. These constructs encode logical dependencies and feasibility constraints directly into the model. Consequently, only plans that satisfy all constraints are generated, and no infeasible configurations are explored. The planning model is divided into a domain description, defining general action structures and system dynamics, and a problem instantiation specifying concrete parameter variants, initial conditions, and optimization objectives. Objective values are computed as:

$$J(\pi) = w_c \cdot C(\pi) - w_q \cdot Q(\pi) \quad (3)$$

where $C(\pi)$ represents accumulated execution and setup costs, and $Q(\pi)$ denotes quality metrics derived from external evaluations.

Quantitative aspects such as cost and quality are incorporated as numerical fluents, enabling weighted objective formulations, where cost is represented as the sum of action-dependent execution and setup costs, and quality is quantified based on deviations of simulation-derived physical indicators (e.g., forces or temperatures) from specified target values. Based on this formulation, the planner solves the constrained optimization problem

$$\pi^* = \arg \min_{\pi \in \Pi_{\text{valid}}} J(\pi) \quad (4)$$

where Π_{valid} denotes the set of all constraint-satisfying plans. As feasibility is enforced directly within the planning model, the optimization is inherently restricted to valid process paths. The architecture remains agnostic with respect to the specific planning algorithm employed, thereby preserving modularity and allowing alternative reasoning mechanisms to be integrated.

The Plan-to-Simulation layer translates symbolic planning results into executable simulation scenarios. The abstract process path generated by the planner is instantiated into a concrete simulation configuration by assigning parameter values, selecting resource

instances, and generating the corresponding simulation scene. Symbolic actions are mapped to executable simulation operations, while parameter variants are translated into configuration settings for the associated assets. In this step, the abstract decision model is converted into a physically interpretable scene configuration. Resource instances are composed and parametrized according to the plan specification, ensuring semantic consistency between the symbolic plan and its physical execution.

The PyBullet Simulation represents the execution layer of the architecture. It initializes, runs, and controls the simulated production process according to the instantiated configuration. Within this environment, kinematic relations, contact interactions, forces, and motion dynamics are approximated to assess the physical plausibility of the generated plan. Simulations are executed under controlled and reproducible conditions. During execution, relevant state variables and interaction signals are recorded, including temporal information, contact events, force profiles, and positional data. These signals serve as quantitative proxies for performance indicators such as feasibility, quality, and time-dependent cost contributions. Simulation therefore complements formal planning by validating execution-level behaviour and exposing physical effects that are not represented in the symbolic model.

Through the Evaluation and Feedback component, the architectural loop is closed. Logged simulation signals are aggregated into decision-relevant performance metrics. Key performance indicators and objective values corresponding to the optimization criteria defined in the planning model are computed, ensuring comparability between symbolic objectives and physically observed outcomes. Based on these results, alternative process paths can be systematically ranked and compared. Furthermore, evaluation outcomes may influence subsequent planning cycles by adjusting objective weightings, introducing additional constraints, or excluding infeasible process parameter variants. This feedback mechanism enables structured reduction of the decision space through successive exploration steps. Owing to its modular design, the architecture can be extended toward multi-line scenarios or more complex resource interactions without altering its fundamental structure, thereby preserving scalability and adaptability.

4. FORMAL MODELING AND PLANNING

Highly variant production systems are characterized by multiple sequential process steps, each offering a set of alternative execution options. For each step $i \in \{1, \dots, n\}$, these alternatives are represented as discrete parameter variants, defining different ways of performing the same abstract operation. Let

$$x_i \in V_i \quad (5)$$

denote the parameter variant selected for process step i , where V_i is the set of admissible variants. Instead of modeling continuous parameter ranges or stochastic variations, the decision space is deliberately discretized. Each parameter variant corresponds to a concrete configuration choice and thus represents an explicit decision variable within a finite, yet combinatorically growing, configuration space. This discretization enables systematic and

formally tractable exploration while maintaining explicit control over combinatorial complexity. The overall decision space is defined as the Cartesian product:

$$D=V1\times V2\times\cdots\times Vn \quad (6)$$

However, not all combinations of parameter variants are feasible. Feasibility is governed by a set of constraints:

$$Ck(x1,\dots,xn),k=1,\dots,m \quad (7)$$

which encode resource capabilities, process dependencies, and admissible parameter combinations. These constraints must be satisfied for a configuration to be valid. Accordingly, the set of feasible configurations is defined as:

$$D_{\text{valid}}=\{x\in D|Ck(x1,\dots,xn)=\text{true}\ \forall k\} \quad (8)$$

The derivation of the decision space follows the Product–Process–Resource–Skill (PPRS) modeling principle. Products, processes, resources, and skills are represented as distinct entities to enable structured reasoning about feasibility and configuration dependencies. Abstract processes describe production activities independently of specific equipment. Resources represent physical entities capable of execution, while skills define the ability of a resource to perform a given process. Concrete tasks emerge through the assignment of parameter variants to skills, thereby forming executable alternatives.

Figure 2 illustrates this modeling principle in an abstracted form for a generic process step, independent of any specific production domain. Discrete parameter variants define alternative execution strategies, which are combined into admissible task instantiations and ultimately form executable process plans.

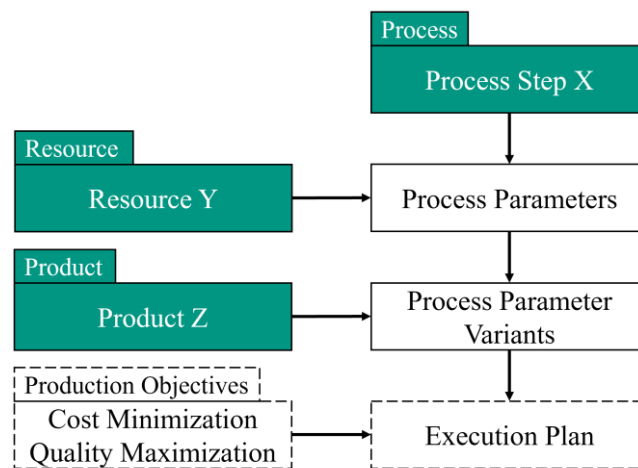


Fig. 2. Abstract representation of the Product–Process–Resource–Skill (PPRS) model used to structure production decisions through discrete parameter variants and resource capabilities

This structured representation is transferable across all process steps by adapting the corresponding parameter sets. The overall decision space is implicitly defined by the combination of admissible parameter variants along the sequential process structure. Rather than enumerating this space explicitly, feasibility constraints and logical dependencies restrict

the set of valid configurations. These constraints arise from resource capabilities, process ordering requirements, and admissible parameter combinations, and are encoded directly in the formal model, thereby defining the feasible solution space D_{valid} .

To enable systematic exploration of this constrained decision space, symbolic planning is employed. Planning problems of this kind are characterized by discrete system states, explicit action choices, and clearly defined goal conditions. The Planning Domain Definition Language (PDDL) is particularly suitable for this purpose, as it supports the unified representation of decision alternatives, constraints, and numerical evaluation criteria. In contrast to gradient-based or continuous optimization approaches, symbolic planning operates directly on discrete, non-convex, and combinatorial decision spaces, where derivative information is unavailable or not meaningful. Furthermore, feasibility is inherently ensured through action preconditions and state transition rules, such that only solutions contained in D_{valid} are generated.

Production objectives are integrated directly into the planning formulation. Typical objectives include cost minimization and quality maximization, where cost is interpreted as an application-specific measure (e.g., setup effort, reconfiguration complexity, or operational overhead), rather than a purely monetary quantity. Quality similarly represents an abstract performance indicator derived from process outcomes. These objectives are formalized as configurable evaluation dimensions and combined within a weighted objective function $J(\pi)$, enabling optimization over the set of feasible process plans.

In this way, formal modeling and symbolic planning provide not only a descriptive representation of process variability, but also an operational mechanism for systematically exploring and evaluating feasible production configurations under competing objective.

5. VERIFICATION SCENARIOS AND RESULTS

To assess the proposed planning-based digital twin architecture, a representative production scenario is considered. The objective of this chapter is not to provide a full validation of production performance, but to verify the functional correctness and internal consistency of the proposed system architecture. Given the large number of parameter variants and the use of exemplary parameter values, the focus lies on demonstrating that the system operates robustly and coherently across the defined decision space. The evaluation therefore emphasizes the traceability of planning decisions, the correctness of generated process paths, and the consistent transfer of planning results into simulation-based execution. At the same time, selected quantitative indicators are evaluated to support this verification and to enable a comparison with baseline approaches.

The central challenge addressed by the proposed architecture is the combinatorial growth of the decision space in highly variant production systems. Each process step may be executed using multiple parameter variants, and each combination of these variants corresponds to a distinct process configuration. When multiple sequential process steps are considered, the resulting number of possible configurations increases rapidly and becomes difficult to analyse manually.

Figure 3 illustrates this structure by representing the decision space as a layered decision graph. Each layer corresponds to a process step, nodes represent selectable parameter variants, and edges indicate feasible transitions between parameter selections of consecutive process steps.

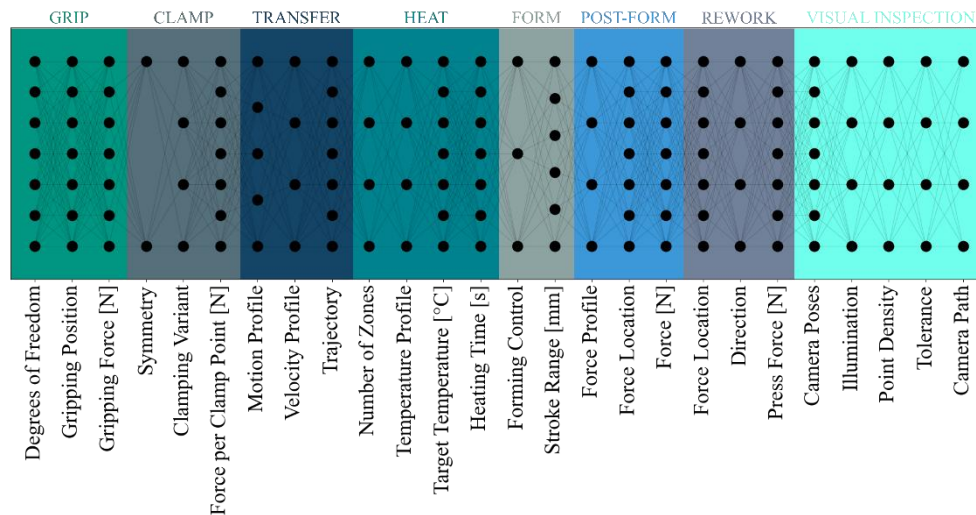


Fig. 3. Physical layout of the thermoforming demonstrator, showing the sequential production stations that define the structure of the planning problem, including gripping, clamping, transfer, heating, forming, post-forming, reworking, and visual inspection

The decision space used in this work is derived from an existing thermoforming demonstrator that serves as the evaluation scenario. The underlying production process and its technological background are described in more detail in [30] and are therefore only summarized here. The demonstrator represents a complete production chain for manufacturing unidirectional tapes and provides a realistic yet controllable environment for evaluating planning-based decision space exploration. Figure 4 shows the physical layout of the demonstrator, including the main processing stations and the overall material flow. The process chain comprises eight sequential steps: gripping, clamping, transfer, heating, forming, post-forming, reworking, and visual inspection. These steps define the structural backbone of the planning problem and form the basis for constructing the decision space shown in Fig. 4.




Fig. 4. Physical layout of the thermoforming demonstrator, showing the sequential production stations that define the structure of the planning problem, including gripping, clamping, transfer, heating, forming, post-forming, reworking, and visual inspection.

Without additional restrictions, the decision graph contains a very large number of possible paths from the initial to the final process step. To demonstrate how the planning approach structures and reduces this space, an exemplary planning scenario is defined in which additional constraints restrict admissible parameter variants. These constraints represent feasibility requirements and production objectives, such as minimum force levels, restricted motion profiles, bounded temperature ranges, or predefined inspection configurations. Instead of enumerating all combinations explicitly, the planner applies these constraints symbolically to eliminate infeasible or undesired branches of the decision graph.

Figure 5 illustrates the resulting reduced decision space for this scenario. Compared to the unconstrained graph shown previously, large parts of the decision space are pruned while feasible and goal-consistent process paths remain accessible. In the investigated scenario, the formalized constraints reduce the theoretical configuration space of about 61.3%. This highlights the effectiveness of constraint-based modelling as a first step in controlling combinatorial complexity before planning is applied.

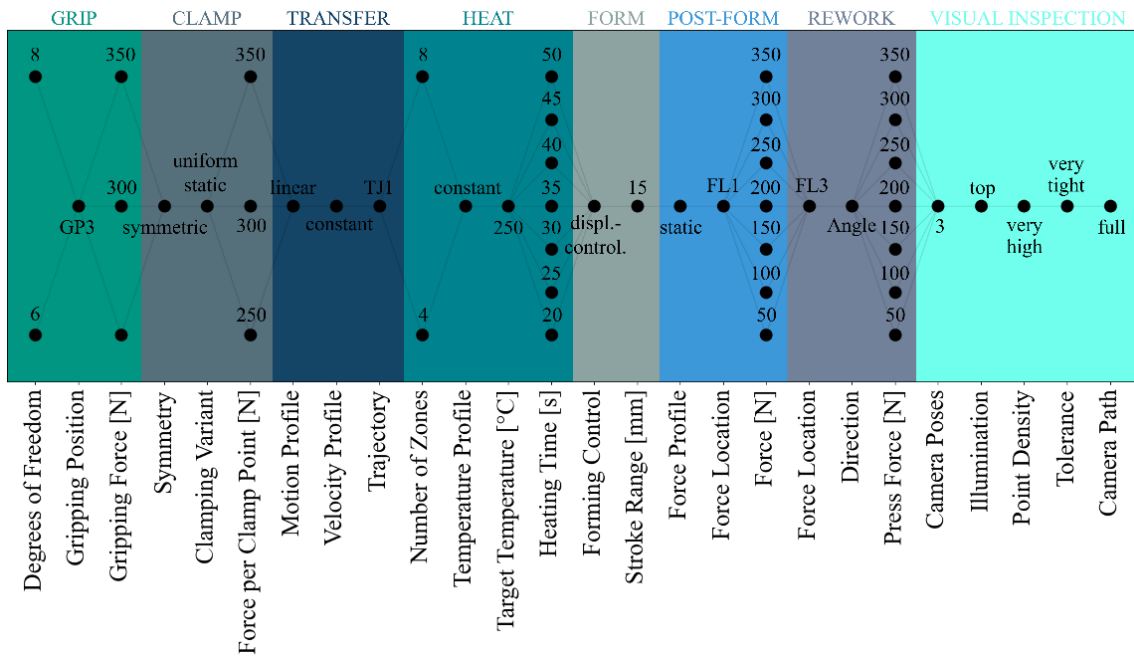


Fig. 5. Reduced decision space resulting from constraint-based planning. Each column represents a process step with discrete parameter variants, while nodes denote admissible selections and edges indicate valid transitions. Infeasible configurations are eliminated through formal constraints, leading to a significantly pruned decision graph that preserves only feasible and goal-consistent process paths

To further assess the effectiveness of the proposed approach, the planning-based method is compared against two baseline strategies: random sampling of parameter combinations and rule-based filtering using simple constraints without planning.

The comparison reveals significant differences in both feasibility and solution quality. Random sampling achieves a feasibility rate of only 2%, indicating that physically valid process configurations are rarely obtained through unstructured exploration. Rule-based filtering improves this rate to approximately 35%, as basic constraints eliminate a portion of infeasible configurations. In contrast, the planning-based approach achieves a feasibility rate of 100%, since feasibility constraints are explicitly encoded in the planning model and enforced during solution generation.

In terms of solution quality, measured by the objective function $J(\pi)$, the planning-based approach consistently identifies the optimal process configuration. While both baselines exhibit a broad distribution of objective values whereas the planning-based approach yields a consistent and optimal result as shown in Fig. 6.

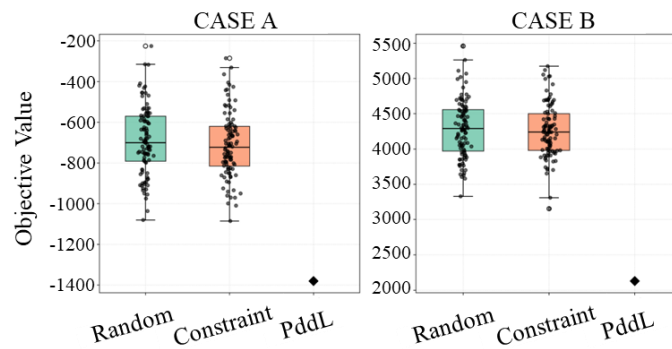


Fig. 6. Distribution of objective values $J(\pi)$ for planning-based and baseline approaches. The planning-based method consistently identifies the optimal solution with minimal variance, while random sampling and rule-based filtering exhibit significantly higher variability and inferior average performance

The feasible process configurations generated by the planner are subsequently transferred to the execution layer of the architecture. Each selected process path defines a concrete sequence of process steps together with the associated parameter selections. In combination with the geometric and physical resource descriptions, these planning results are used to instantiate executable simulation scenarios.

Figure 6 depicts an exemplary state of the resulting PyBullet simulation environment. The simulation scene is automatically generated from the planning results by instantiating the required resources, assigning the selected parameter configurations, and executing the corresponding process sequence. At this stage, the simulation primarily serves to verify the correct execution of the generated plans and the consistency of the plan-to-simulation transfer, rather than to evaluate absolute production performance.




Fig. 7. Example PyBullet simulation scene generated from a planned process configuration for execution-level verification

Overall, the presented results demonstrate that the proposed architecture enables consistent and traceable exploration of large decision spaces in highly variant production systems. The verification confirms that formal modelling and constraint-based planning reliably generate coherent process paths and that these planning results can be transferred

consistently into simulation-based execution. At the same time, the quantitative comparison shows that planning not only guarantees feasibility but also systematically identifies superior solutions compared to baseline approaches. The results therefore substantiate both the functional correctness of the architecture and the effectiveness of planning-based decision space exploration within the proposed digital twin framework.

6. CONCLUSION & OUTLOOK

This paper presented a planning-based digital twin architecture for structured decision space exploration in highly variant and immature production processes. The core contribution lies in the integration of formal decision space modeling, symbolic planning, physics-based simulation, and evaluation-driven feedback into a unified and modular architecture. Rather than addressing individual process parameters in isolation, the proposed approach enables systematic reasoning over combinatorial configuration spaces and supports the generation, execution, and assessment of consistent process alternatives.

The architecture combines a discretized decision space representation with a symbolic planning layer and a PyBullet-based execution environment. By structuring process variability through discrete process parameter variants and explicitly modeling feasibility constraints and production objectives, large decision spaces can be explored in a controlled and traceable manner. The verification scenarios demonstrated that the proposed pipeline reliably generates coherent process sequences, applies constraints consistently, and transfers planning results into executable simulation scenarios. These results confirm the functional correctness of the architecture and its suitability for decision space structuring and exploration in highly variant production settings.

Beyond functional verification, the quantitative evaluation further highlights the effectiveness of the proposed approach. The results show that planning-based decision generation significantly outperforms baseline methods in both feasibility and solution quality. While random sampling achieves a feasibility rate of only 2% and rule-based filtering improves this to approximately 35%, the planning-based approach consistently produces feasible solutions with a rate of 100%. Furthermore, the quality of generated solutions, measured by the objective function $J(\pi)$, is substantially improved. The planner identifies an optimal solution, whereas baseline methods exhibit considerably higher average objective values and large variability. These findings underline that planning not only ensures constraint satisfaction but also enables systematic identification of superior process configurations.

A key outcome of this work is the demonstration that planning-based methods provide an effective mechanism for managing combinatorial complexity in early-stage or exploratory production processes. Instead of exhaustive enumeration, the decision space is systematically reduced through formal constraints and objective-driven planning. In this sense, the digital twin is not limited to passive representation or simulation, but actively supports decision-making by coupling symbolic reasoning with execution-level verification.

While the presented results verify the functional behavior of the proposed architecture, several directions for future work remain. A central next step is the coupling of the planning

and simulation pipeline with a fully realized physical demonstrator. Such a setup would enable experimental validation under real process conditions and allow bidirectional data exchange between the physical system and the digital twin. Furthermore, the integration of higher-fidelity process models, for example through FEM-based material simulations, would allow more accurate representation of physical effects in forming operations.

To maintain computational efficiency, surrogate models could be introduced to approximate detailed simulations and enable scalable evaluation across larger decision spaces. This would support more extensive exploration while preserving physical relevance. In addition, the architecture is inherently extensible toward more complex production scenarios, including multi-line systems and parallel resource utilization. Incorporating such scenarios would require additional decision dimensions, such as line assignment or resource sharing, which can be naturally integrated into the existing planning framework.

Overall, the proposed architecture provides a robust foundation for planning-based digital twins in highly variant production environments. By combining formal decision space modeling, symbolic planning, and simulation-based execution within a unified framework, it offers a scalable and extensible approach for systematic decision support in early-stage and evolving production systems.

ACKNOWLEDGEMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project 5339

REFERENCES

- [1] Cluster Elektromobilität Süd-West. Die Produktionskette “Batteriesystem” und kritische Ressourcen. https://www.emobil-sw.de/fileadmin/media/emobilbw/Publikationen/Studien/Themenpapier_Cluster_ESW_Produktionskette_Batteriesystem_und_kritische_Ressourcen.pdf. Accessed 19 April 2026.
- [2] EYDLIN A., 2022, *EU-Kommission Billigt Staatshilfen Für Zweites Batteriegroßprojekt*. <https://www.zeit.de/mobilitaet/2021-01/elektroauto-eukommission-batterie-foerderung-tesla>.
- [3] WESTKÄMPER E., LÖFFLER C., 2016, *Strategien Der Produktion: Technologien, Konzepte Und Wege In Die Praxis*. Springerlink Bücher, Berlin, Heidelberg: Springer Vieweg, 309.
- [4] Storz, T., Schild, L., Vogt, M., Hofmann, J., Fleischer, J., & Lanza, G. (2021). Entwicklung eines flexiblen Produktionssystems für Batteriezellen. *Zeitschrift für wirtschaftlichen Fabrikbetrieb*, 116(4), 247-253.
- [5] KERN W., HRSG. 1996, *Handwörterbuch Der Produktionswirtschaft*. 2., Völlig Neu Gestaltete Aufl, Enzyklopädie der Betriebswirtschaftslehre 7, Stuttgart: Schäffer-Poeschel, 2462.
- [6] WATERMEYER K., 2016, *Ablaufplanung Mit Alternativen Prozessplänen*, 1, Aufl. 2016, Best-Masters, Wiesbaden: Springer Gabler, 90 S.
- [7] YUSUF Y.Y., SARHADI M., GUNASEKARAN A., 1999, *Agile Manufacturing*, *International Journal of Production Economics*, 62/12, 33–43.
- [8] OVERBECK L., BRÜTZEL O., STRICKER N., LANZA G., 2020, *Digitaler Zwilling Des Produktionssystems*, *Zeitschrift für wirtschaftlichen Fabrikbetrieb*, 115, 62.
- [9] FULLER A., FAN Z., DAY C., BARLOW C., 2020, *Digital Twin: Enabling Technologies, Challenges and Open Research*, *IEEE Access* 8. Conference Name: IEEE Access, 108952–108971.
- [10] GRIEVES M., 2014, *Digital Twin: Manufacturing Excellence Through Virtual Factory Replication*, Whitepaper.

- [11] TAO F., ZHANG M., 2017, *Digital Twin Shop-Floor: a New Shop-Floor Paradigm Towards Smart Manufacturing*, IEEE Access 5, Conference Name: IEEE Access, 20418–20427.
- [12] LENG J., WANG D., SHEN W., LI X., LIU Q., CHEN X., 2021, *Digital Twins-Based Smart Manufacturing System Design in Industry 4.0: A Review*, Journal of Manufacturing Systems 60, 119–137.
- [13] MIHAI S., YAQOUB M., HUNG D.V., DAVIS W., TOWAKEL P., RAZA M., KARAMANOGLU M., BARN B., SHETVE D., PRASAD R.V., VENKATARAMAN H., TRESTIAN R., NGUYEN H.X., 2022, *Digital Twins: a Survey on Enabling Technologies, Challenges, Trends and Future Prospects*, IEEE Communications Surveys & Tutorials 24.4. Conference Name: IEEE Communications Surveys & Tutorials, S. 2255–2291.
- [14] ERAMO R., et al., 2022, *Conceptualizing Digital Twins*, IEEE Software 39/2, 39–46, <https://doi.org/10.1109/MS.2021.3130755>.
- [15] MOURTZIS D., DOUKAS M., BERNIDAK, D., 2014, *Simulation in Manufacturing: Review and Challenges*, Procedia CIRP 25, 213–229, <https://doi.org/10.1016/j.procir.2014.10.032>.
- [16] NEGRI E., FUMAGALLI L., MACCHI M., 2017, *A Review of the Roles of Digital Twin in CPS-Based Production Systems*, Procedia Manufacturing 11, 939–948, <https://doi.org/10.1016/j.promfg.2017.07.198>.
- [17] LIU C.K., NEGRUT D., 2021, *The Role of Physics-Based Simulators in Robotics*, Annual Review of Control, Robotics, and Autonomous Systems 4, 35–58, <https://doi.org/10.1146/annurev-control-072220-093055>.
- [18] COUMANS E., BAI Y., 2016, *Pybullet Quickstart Guide*.
- [19] KOENIG N., HOWARD A., 2004, *Design and Use Paradigms for Gazebo, an Open-Source Multi-Robot Simulator*, 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Bd, 3, 2149–2154, 3.
- [20] Baumgärtner, J., Hansjosten, M., Schönhofen, D., & Fleischer, I. J. (2023). Pybullet industrial: A process-aware robot simulation. *Journal of Open Source Software*, 8(85), 5174. FROMMER J., SCHLEIPEN M., BEYERER J., 2013, *PPRS: Production Skills and Their Relation to Product, Process, and Resource*, 2013 IEEE 18th Conference on Emerging Technologies & Factory Automation (ETF A). 2013 IEEE 18th Conference on Emerging Technologies & Factory Automation (ETF A), ISSN: 1946–0759, 1–4.
- [21] GHALLAB M., NAU D., TRAVERSO P., 2004, *Automated Planning: Theory and Practice*, Amsterdam, Morgan Kaufmann, 635 S.
- [22] GHALLAB M., HOWE A., KNOBLOCK C., MCDERMOTT D., RAM A., VELOSO M., WELD D., WILKINS D., BARRETT A., CHRISTIANSON D., FRIEDMAN M., KWOK C., GOLDEN K., PENBERTHY S., SMITH D.E., SUN Y., 1998, Pddl - the Planning Domain Definition Language.
- [23] AIELLO M., GEORGIEVSKI I., 2024, *Introduction to AI Planning*, arXiv, 2412, 11642[cs].
- [24] PUROHIT L., CHOUHAN S.S., 2023, *Web Service Composition Using an AI Planning Technique*, Innovative Engineering with AI Applications, John Wiley & Sons, Ltd, S, 65–82.
- [25] GRIEVES M., VICKERS J., 2016, *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behaviour in Complex Systems*, Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches. Hrsg. von F.-J. Kahlen; S. Flumerfelt & A. Alves, Cham: Springer International Publishing, 85–113.
- [26] TAO F., ZHANG H., LIU A., NEE A.Y.C., 2018, *Digital Twin in Industry: State-of-the-Art*, IEEE Transactions on Industrial Informatics 15.4. Conference Name: IEEE Transactions on Industrial Informatics, 2405–2415.
- [27] KRITZINGER W., KARNER M., TRAAR G., HENJES J., SIHN W., 2018, *Digital Twin in Manufacturing: a Categorical Literature Review and Classification*, IFAC-PapersOnLine, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018 51.11, 1016–1022.
- [28] JOPPICH T.D., 2019, *Beitrag Zum Umformverhalten Von PA6/CF Gelegelaminaten Im Nichtisothermen Stempelumformprozess*, Diss. Karlsruhe: Karlsruher Institut für Technologie.