SMART PROCESS PLANNING FOR AUTOMATED FIBER PLACEMENT

A Dissertation
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
Automotive Engineering

by
Matthew John Godbold
August 2026

Accepted by:
Dr. Ramy Harik, Committee Chair
Dr. Saeed Farahani
Dr. Robert Prucka
Dr. Vinita Jansari

PLAIN LANGUAGE ABSTRACT

Many industries, including aerospace, automotive, wind energy, maritime, and sporting goods, rely on strong, lightweight materials called composites. These materials are made by layering fibers, which can come in the form of narrow strips or wider sheets, and setting them in a polymer matrix. One of the most advanced ways to make these parts is through automated fiber placement, where a machine lays down the fibers in precise patterns. This method can create very efficient and strong designs, but it is complex, expensive, and often depends heavily on the experience of skilled engineers.

Today, the design, manufacturing, and inspection stages of composite production are usually handled separately. This separation means that important information, such as how a part will be built or what defects might occur, is not always shared between stages. As a result, parts may not be as lightweight, strong, or defect-free as possible, and the process can take longer and cost more.

This research develops a smart process planning system that connects design, manufacturing, and inspection into one continuous process. Built as software that works with existing tools, the system can automatically plan how the fibers are placed, predicting and reducing defects while improving both manufacturability and strength. The system optimizes not only individual layers but also how defects are distributed across all layers, preventing them from stacking up in ways that weaken the final part. It also uses inspection results from completed parts to improve future designs, creating a feedback loop where each stage informs the others.

The system was tested by designing a composite panel using this new approach and comparing it to a panel made with state-of-the-art manual planning methods. The results showed that the system could intentionally control where defects appeared and increase the efficiency of the planning process.

By unifying design, manufacturing, and inspection, this research shows a way to make advanced composite manufacturing more efficient, consistent, and cost-effective. This approach lowers the barrier to using automated fiber placement and opens the door for its wider adoption not only in aerospace but also in industries such as automotive, wind energy, maritime, and sporting goods, where strong and lightweight structures are essential.

ABSTRACT

Automated Fiber Placement (AFP) is a preeminent manufacturing technology for producing high-performance composite structures, offering precise material placement, reduced waste, and the potential for highly optimized lightweight designs. While AFP has been widely adopted in aerospace applications, its broader use across industries such as automotive, wind energy, maritime, and sporting goods remains limited. A key barrier to wider adoption is the traditional separation of design, manufacturing, and inspection processes, which drives up cost and results in inefficiencies, suboptimal designs, and limited feedback between lifecycle stages. In current practice, manufacturing is treated as an isolated process, disconnected from early-stage design decisions and post-production inspection data. This siloed approach attenuates the potential of AFP to deliver fully optimized, defect-minimized composite structures.

This dissertation addresses these limitations by developing the Integrated Smart Process Planning (iSPP) framework for AFP, which supersedes the traditional manual, *a posteriori* approach with a closed-loop, data-driven process. While current AFP process planning is heavily dependent on tacit knowledge and offers little systematic feedback between stages, the iSPP framework transforms this workflow into a smart, adaptive, and interconnected process. It guides manufacturing decisions by reconciling original design intent with real-world production data.

The research is realized through the development of a set of integrated capabilities that collectively form the iSPP framework. Each capability was designed to address a specific gap in the AFP lifecycle while interfacing seamlessly to create a cohesive, end-to-end process. At the core is the Computer-Aided Process Planning (CAPP) software, which semi-automates process planning decisions and manages ply- and laminate-level data to enable informed trade-offs between manufacturing feasibility and design intent. Building on

this foundation, novel defect scoring algorithms were developed to quantify manufacturing quality in a way that supports optimization and decision-making. These feed into ply-level optimization methods that combine global surveying with a local perpendicular shift approach to minimize in-plane defects. For laminate-level optimization, two complementary approaches are presented: a stagger-based method that systematically distributes starting points to reduce through-thickness defect alignment, and a gravity-based model derived from an analogy to Newton's law of universal gravitation that quantifies spatial defect interactions to minimize through-thickness stacking while protecting margin-critical areas. To bridge design and manufacturing, the platform integrates with Collier Aerospace's HyperX to incorporate structural margin-of-safety data directly into defect scoring, and with Vericut Composite Programming (VCP) to automate ply generation, perform defect prediction, and produce NC programs. Finally, the platform incorporates inspection data to compare as-manufactured and as-designed panels. The full capability is demonstrated through verification and demonstration on the Quartic benchmark surface, a doubly-curved saddle geometry developed by Collier Aerospace as a challenging AFP process planning test case. The CAPP-driven approach was compared to manual planning baselines, demonstrating intentional control over defect placement, improved process planning efficiency, and the potential for structural performance benefits through margin-informed optimization.

The results demonstrate that integrating design, manufacturing, and inspection into a unified AFP process planning platform enables intentional, data-driven control over build quality and defect distribution. By embedding defect awareness, structural performance considerations, and inspection feedback into a single, adaptive environment, the iSPP framework reduces reliance on empirical methods and lowers the barrier to adoption across diverse industries. This lifecycle-integrated approach provides a foundation for next-generation AFP systems that are more accessible, adaptable, and capable of delivering high-quality composite structures across a broad range of manufacturing domains.

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere gratitude to my advisor, Dr. Ramy Harik, for his unwavering support, mentorship, and encouragement throughout this journey. His guidance has been instrumental in shaping both this research and the person I have become. I am also deeply grateful to my committee members, Dr. Saeed Farahani, Dr. Robert Prucka, and Dr. Vinita Jansari, whose insights, feedback, and thoughtful questions have strengthened this work considerably. I would also like to extend special thanks to Dr. Brian Tatting, whose expertise and guidance have been instrumental in the development of this research.

I owe a debt of gratitude to my colleagues and fellow researchers who have contributed to this work in countless ways. Your collaboration, camaraderie, and willingness to share ideas have made this endeavor both possible and enjoyable. My warmest thanks go to the members of the Clemson Composites Center and Clemson University for their friendship and steadfast support.

The material contained in this dissertation is based upon work supported by NASA under award **No. 80NSSC22M0155**. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Aeronautics and Space Administration.

Finally, I must express my profound appreciation to my family and friends. Their unconditional love, patience, and encouragement throughout my academic journey have been the foundation upon which all of this was built. I am who I am today because of their unwavering belief in me.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

The use of composite materials has marked a significant turning point in engineering innovation, enabling the design of structures that are lighter, stronger, and more durable than those made from traditional metals. The concept of combining distinct materials to achieve superior properties is not new. Ancient civilizations mixed straw with mud to form reinforced bricks [1], while Mongolian warriors crafted bows from layers of wood, horn, and sinew to achieve superior performance [2]. These early examples demonstrated the fundamental principle that combining different constituents can yield properties unattainable by any single material alone.

In the modern era, the development of advanced composites has transformed industries such as aerospace, automotive, maritime, wind energy, and sporting goods. By embedding high-strength fibers such as carbon or glass within a polymer matrix, engineers can create materials with exceptional strength-to-weight and stiffness-to-weight ratios, corrosion resistance, and design flexibility [3, 4]. These attributes have enabled lighter, more efficient structures, directly translating into reduced fuel consumption, lower emissions, and improved performance. Today, advanced composites make up more than half the structural weight of state-of-the-art aircraft such as the Boeing 787 and Airbus A350 [5, 6].

The demand for larger, more complex, and higher-performance composite structures has driven the adoption of automated manufacturing technologies. Among these, AFP has emerged as a leading process, capable of precisely laying down narrow strips of fiber in controlled orientations to build complex geometries with minimal waste [7]. AFP offers the promise of high repeatability, reduced labor, and the ability to tailor material placement to structural requirements. Its adoption has been particularly strong in aerospace, where reducing weight directly translates to improved fuel efficiency and performance [8]. Other industries have also begun to explore AFP, with the automotive sector investigating lightweight body panels and suspension components, and the wind energy sector evaluating its use in producing large turbine blades that must withstand harsh operating conditions [9, 10].

Despite these advantages, AFP has not yet reached its full potential. The process remains costly and complex, often relying on trial-and-error methods and the tacit knowledge of experienced engineers. More critically, the lifecycle of composite manufacturing—spanning design, process planning, manufacturing, and inspection—remains fragmented. Each stage is typically treated in isolation, limiting feedback, increasing inefficiencies, and driving up cost. Consequently, AFP remains underutilized outside of aerospace, where the high cost can be justified, and adoption in industries such as automotive, maritime, and wind energy has been limited.

This dissertation addresses these challenges by developing the iSPP framework, a smart, lifecycle-integrated approach to AFP. By unifying design, manufacturing, and inspection into a closed-loop system, the research aims to reduce costs, improve efficiency, and enable intentional control over defects and structural performance. This vision builds on the concept of "Smart Automated Fiber Placement," where AFP evolves into an interconnected, data-driven process that leverages digital tools, feedback integration, and adaptive decision-making to improve quality and accessibility across industries [11].

The significance of composites in modern engineering cannot be fully appreciated without first understanding their fundamental makeup and evolution. While the preceding discussion has highlighted their historical context and growing importance across industries, a more detailed treatment is necessary to establish the foundation for understanding why processes such as AFP have become essential for producing lightweight, high-performance structures. This chapter provides that foundation and motivates the research presented in subsequent chapters. Section 1.1 introduces composite materials, their constituents, and the manufacturing methods that have shaped their evolution. Section 1.2 presents AFP in detail, covering process fundamentals, machine types, key process parameters, and the technology's historical development. Section 1.3 identifies the shortcomings of current AFP practice that limit its broader adoption. Section 1.4 introduces the vision of Smart AFP and positions the iSPP framework as a response to these challenges. Finally, Section 1.5 summarizes the dissertation's contributions and outlines the structure of the remaining chapters.

1.1 OVERVIEW OF COMPOSITE MATERIALS AND MANUFACTURING

Composite materials are engineered systems composed of two or more distinct constituents that, when combined, exhibit properties superior to those of the individual components. The reinforcement phase typically consists of high-strength, high-stiffness fibers such as carbon, glass, or aramid. The matrix phase—usually a polymer, though metals and ceramics are also used—binds the fibers, transfers loads between them, and provides environmental protection [4]. The resulting material system offers high specific strength and stiffness, corrosion resistance, and the ability to tailor properties through fiber orientation and stacking sequence. These attributes have made composites indispensable across a wide range of industries, from aerospace to sporting equipment [12].

The structural form most commonly used in high-performance applications is the laminated composite. Laminates are created by stacking multiple thin plies, each containing unidirectional or woven fibers embedded in a matrix, at prescribed orientations. By varying the fiber angles, thickness, and stacking sequence, engineers can tailor the stiffness and strength of the laminate to meet specific load requirements [13]. This hierarchical nature of composites, where material behavior is governed at multiple scales from fiber-matrix interactions to full structural performance, is illustrated in Figure 1.1.

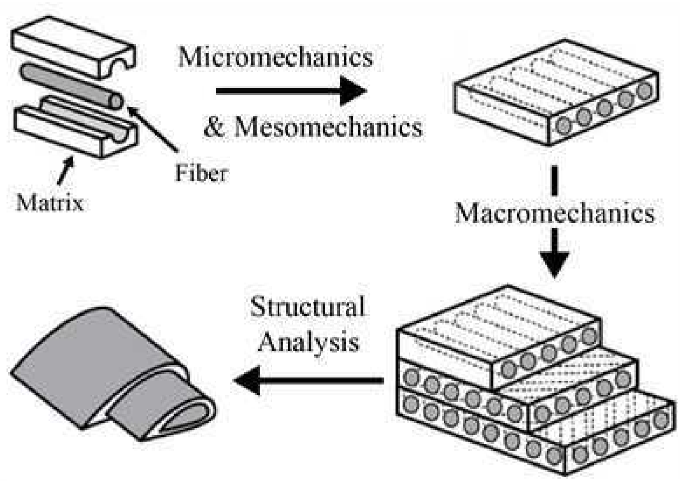

Figure 1.1: Hierarchical representation of composite materials, from fiber and matrix interactions (micromechanics and mesomechanics) to laminate behavior (macromechanics) and structural performance (reproduced from [14]).

1.1.1 MANUFACTURING METHODS

The manufacturing of high-performance composite parts, particularly in aerospace applications, has historically relied on manual layup processes. Manual layup requires meticulous cutting and placement of unidirectional or woven fiber mats onto an open mold, following the stacking sequence until the desired part thickness is achieved. These mats may contain dry fibers or pre-impregnated fibers, commonly referred to as prepreg. In the dry process, fibers are placed on the mold surface and resin is either poured or infused under vacuum using techniques such as Resin Transfer Molding (RTM) or Vacuum Assisted Resin Transfer Molding (VARTM). To avoid the risk of improper resin transfer, high void content, or low fiber volume fraction, many applications use prepreg materials.

The hand layup of prepreg mats mirrors that of dry mats, with an optional debulking process between layers to minimize air entrapment. Curing generally occurs under high pressure and temperature in an autoclave or oven, further reducing voids. To mitigate the high capital and operational costs of autoclaves, significant research has focused on developing Out-of-Autoclave (OoA) prepregs, which cure in an oven under vacuum pressure alone.

For more complex components, flat tools may be used initially, with a subsequent forming process applied to achieve the final shape. While effective, manual layup is labor-intensive, prone to human error, and often results in defects due to surface complexity. Quality control relies heavily on costly manual inspection, although more recent innovations have sought to implement automated inspection systems. Despite its continued use, the throughput and accuracy of hand layup fall short of the demands of modern manufacturing, highlighting the need for automated techniques.

Before the emergence of AFP, large composite structures were manufactured using ATL and FW, processes that remain in use for certain applications today. These methods are depicted in Figure 1.2. ATL (Figure 1.2a), introduced in the 1970s, employs wide prepreg tapes with a typical width of 3 to 12 in. and is efficient for flat or mildly curved surfaces. It automates backing removal and allows variation of process parameters such as temperature, compaction, and tape tension. FW (Figure 1.2b) is primarily used to produce hollow or cylindrical components such as pipes and tanks by winding continuous fiber around a rotating mandrel in a predetermined pattern.

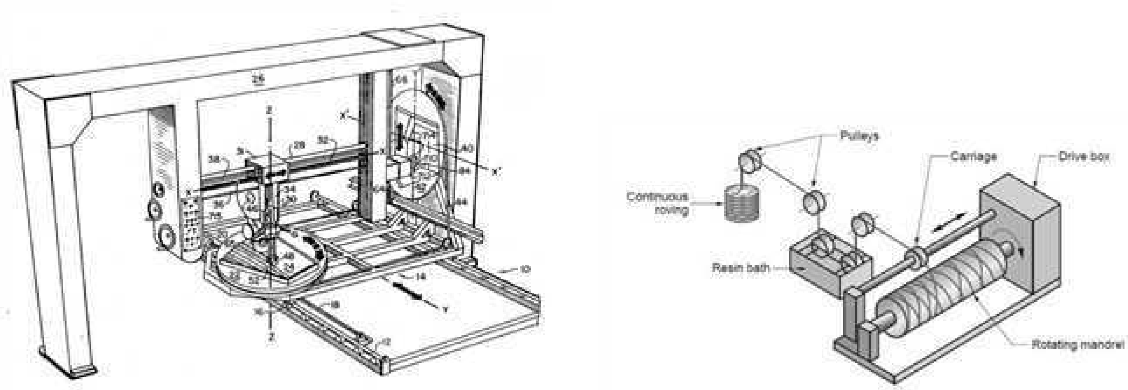

(a) (b)

Figure 1.2: (a) ATL [15] and (b) FW [16].

The transition from ATL to AFP began when engineers introduced mechanisms to slit wide tapes into narrow strips, or tows. In 1974, Hardesty et al. patented an ATL head capable of slitting tape into strips approximately 0.125 in. wide [17], enabling placement on more complex geometries than were practical with wider tapes. AFP builds on this foundation by replacing wide tapes with multiple individually controlled tows, typically 0.25 to 1.0 in. width. This modification reduces material waste, increases process flexibility, and allows the production of highly contoured surfaces. The following section examines AFP in detail, covering its process fundamentals, machine configurations, key parameters, and historical evolution.

1.2 AUTOMATED FIBER PLACEMENT (AFP)

AFP is an advanced composite manufacturing process that enables the precise placement of continuous fiber tows to form multi-layered laminates. It is considered one of the most versatile and capable automated methods for producing large, high-performance composite structures, particularly in aerospace applications.

1.2.1 PROCESS DESCRIPTION

In AFP, a robotic or gantry-based system is equipped with a fiber placement head that delivers multiple narrow tows onto a tool surface. These tows, typically 0.25 to 1.0 in. width, are stored on spools housed either in a creel or directly on the placement head. The AFP head contains all elements needed to deliver the fiber, heat the underlying layer to provide tackiness, and compress the incoming material onto the substrate [18]. A heat source at the nip point—the location where the compaction roller contacts the tool surface—activates the resin to ensure adhesion between the incoming tows and the underlying surface. Figure 1.3 shows an example of a robotic AFP machine: the Integrated Structural Assembly of Advanced Composites (ISAAC) system developed by Electroimpact and located at NASA Langley Research Center.

A single pass of the AFP head deposits a group of tows, known as a course. Multiple courses form a ply, and multiple plies stacked in sequence form a laminate. By controlling tow placement, cut-and-restart operations, and steering, AFP can produce laminates with complex geometries and tailored fiber orientations. This flexibility allows AFP to minimize material waste, reduce defects, and achieve high repeatability compared to manual layup or ATL [18, 19].

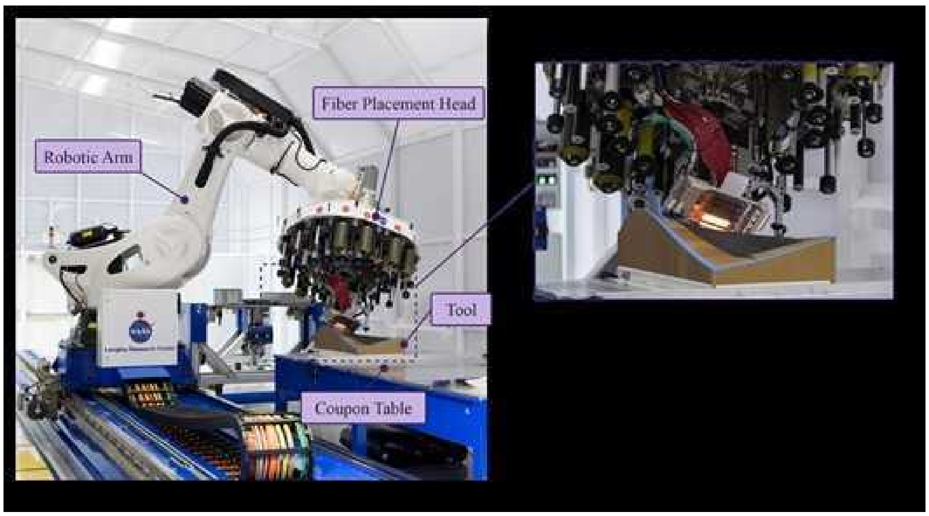

Figure 1.3: Fabrication of manufacturing demonstration unit [20].

1.2.2 MACHINE TYPES

The evolution of AFP machines reflects their origins as a hybrid between FW and ATL, with advancements that have significantly improved their versatility and precision. These machines are typically categorized into three primary types: vertical gantry, horizontal gantry, and robotic systems, each with distinct capabilities suited to different manufacturing needs.

1.2.2.1 VERTICAL GANTRY SYSTEMS

Vertical gantry AFP systems are among the most commonly employed in aerospace manufacturing, primarily due to their ability to handle large structures such as airplane wings and turbine blades. These systems operate on a Cartesian frame, allowing the gantry to move along predefined axes to deliver material across the length of the structure. The AFP head, mounted on the gantry, is generally equipped with three degrees of freedom, enabling it to place material on flat, concave, convex, or even doubly curved surfaces.

As illustrated in Figure 1.4, a typical vertical gantry AFP system features a large working envelope suitable for high-precision layup of extensive composite parts. These systems are often equipped with refrigerated creels to prolong the out-life of thermoset materials, preventing premature curing and reducing the incidence of defects such as fuzz generation and resin buildup.

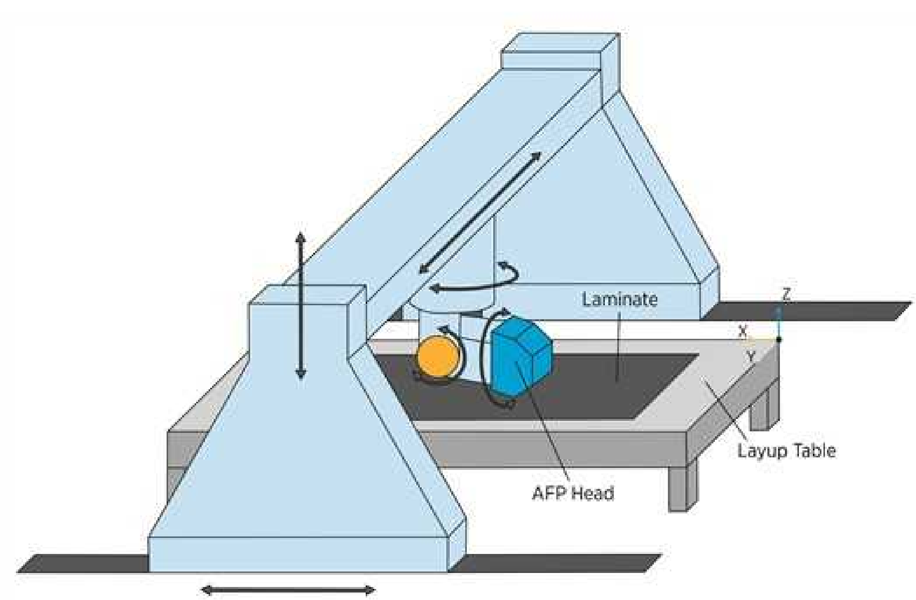

Figure 1.4: Vertical gantry AFP machine [18].

1.2.2.2 HORIZONTAL GANTRY SYSTEMS

Horizontal gantry AFP systems share similarities with their vertical counterparts but differ in their orientation. In these systems, the AFP head travels horizontally along the gantry, making them particularly effective for manufacturing cylindrical structures such as fuselage barrel sections. Figure 1.5 illustrates a horizontal gantry configuration, representative of those used to produce the Boeing 787 fuselage barrels. While less common in recent years than vertical systems, horizontal gantry AFP machines offer unique advantages, including the ability to lay up material adjacent to the structure without encroaching on the space needed for layup.

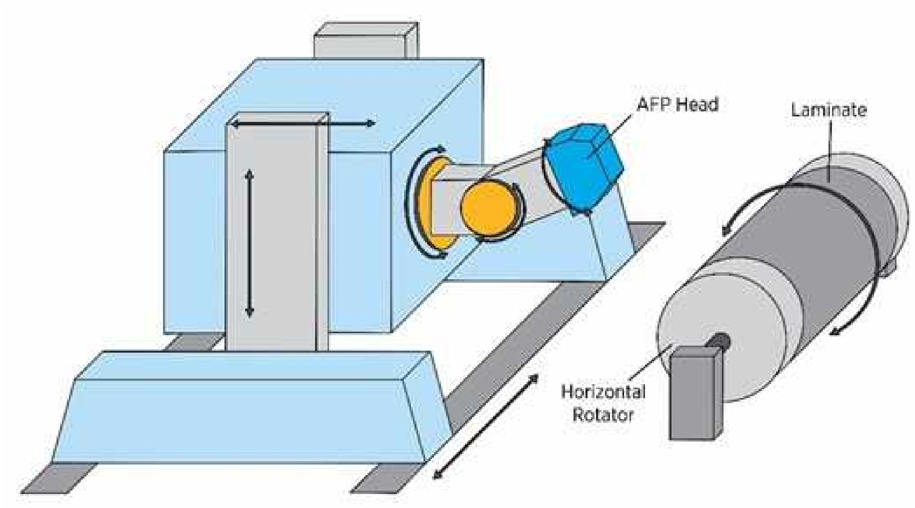

Figure 1.5: Horizontal rotator AFP machine [18].

1.2.2.3 ROBOTIC SYSTEMS

Robotic AFP systems represent the latest innovation in fiber placement technology. Unlike gantry systems, where material delivery is primarily managed by the movement of the gantry, robotic systems rely on the mobility of the robotic arm. These systems are characterized by their versatility, as the robotic arm provides multiple degrees of freedom, enabling the AFP head to access intricate geometries with ease.

As depicted in Figure 1.6, robotic AFP systems can be configured with a stationary layup table (Figure 1.6a) or as collaborative robotic systems (Figure 1.6b) that enhance adaptability for research and manufacturing environments. These systems are ideal for both research and industrial applications, offering the flexibility to swap heads for different processes, such as stitching or projection. Recent advancements have also introduced dual robotic systems and alternative configurations where the tool is manipulated by the robotic arm, further enhancing the flexibility of the AFP process.

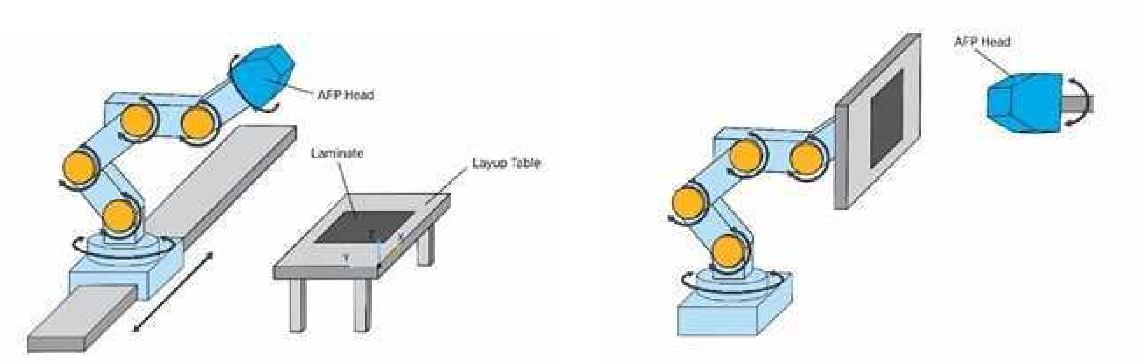

(a) (b)

Figure 1.6: (a) Robotic AFP with layup table and (b) collaborative robotic AFP system [18].

1.2.3 MATERIALS

The selection of materials in AFP plays a crucial role in determining both the manufacturing process and the quality of the final composite structure. The materials used in AFP are generally categorized into three main types: thermoset, thermoplastic, and dry fiber. Each of these materials presents unique characteristics that influence the processing techniques, operational windows, and the performance of the product.

1.2.3.1 THERMOSET MATERIALS

Thermoset prepregs are the most widely used materials in AFP manufacturing, largely due to their ease of processing. These materials consist of fibers impregnated with a thermoset resin, a polymer that cures to form a rigid, three-dimensional lattice structure. One of the primary advantages of thermosets is that the deposition process occurs at or near room temperature, which simplifies the heating requirements for the machine head and allows the material's natural tack to adhere to the tool surface. However, thermoset materials necessitate a secondary processing step, such as autoclave consolidation, to fully cure the resin at elevated temperatures, which can increase both processing time and cost. Careful temperature control is essential during the AFP process, as temperatures above approximately 70 °C can prematurely initiate the curing reaction, leading to defects. To maintain their workability, thermoset prepregs are stored in a frozen state to slow down the curing process until they are ready to be used.

1.2.3.2 THERMOPLASTIC MATERIALS

Thermoplastic prepregs differ significantly from thermosets in both composition and behavior. The fibers in thermoplastic materials are impregnated with a thermoplastic, which does not cure upon heating but instead can repeatedly transition between solid and fluid states. This recyclability and rework ability are among the key advantages of thermoplastics, along with their high-temperature performance, impact resistance, and extended shelf life at room temperature. Additionally, thermoplastic AFP processes offer the potential for in-situ consolidation, eliminating the need for a secondary curing step, thus reducing overall production time. However, these materials present challenges during layup due to their higher required processing temperatures, typically around 400 °C, and a narrower operational window, which necessitates precise control of process parameters to avoid defects.

1.2.3.3 DRY FIBERS

Dry fibers, as the name suggests, are not pre-impregnated with resin, which sets them apart from both thermoset and thermoplastic prepregs. The absence of resin offers several benefits during the AFP process, such as reduced resin buildup in the machine head, leading to longer maintenance intervals and improved reliability. Dry fibers also allow for easier steering, as the lack of a resin matrix provides greater flexibility for the fibers to bend or shear during layup.

However, the absence of tackiness in dry fibers requires the use of a binder, typically a heat-activated thermoplastic powder or filament that provides the necessary adhesion to hold the tows in place during layup. Dry fiber materials also require a post-layup resin infusion step to achieve the desired mechanical properties, adding complexity to the manufacturing process. This trade-off, however, enables OoA processing: the ability to infuse and cure parts without an autoclave yields reductions in both capital equipment costs and part size constraints, while also improving production rates [21].

1.2.4 PROCESS PARAMETERS

The success of AFP relies heavily on the careful control of several key process parameters, which directly influence the quality of the layup and the mechanical properties of the final composite part. These parameters include heating, speed, compaction pressure, and tow tension, each of which must be precisely managed to achieve optimal results. A simplified schematic of the AFP head, shown in Figure 1.7, illustrates the primary components responsible for applying these parameters during the layup process.

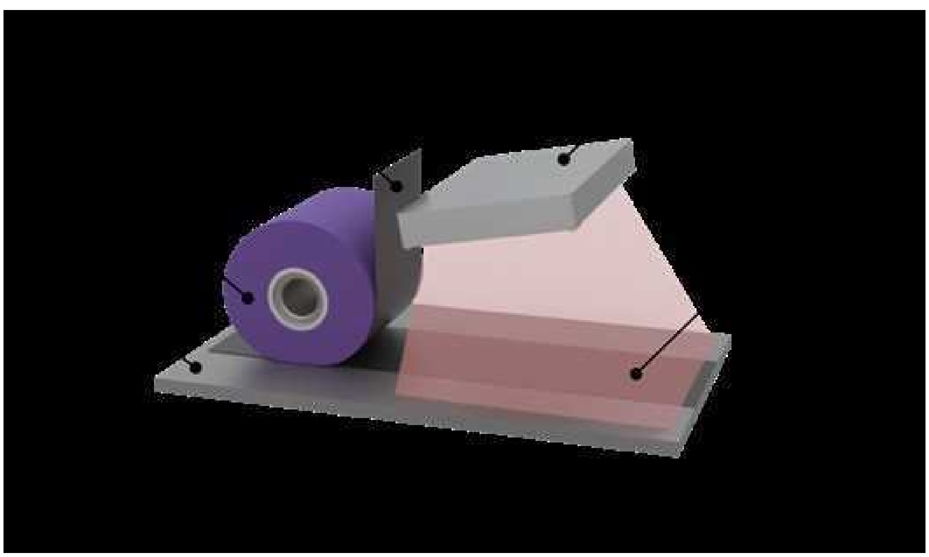

Figure 1.7: Simplified AFP schematic illustrating the key components of the fiber placement head [22].

1.2.4.1 HEATING

Heating is a critical parameter in AFP, essential for ensuring proper adhesion between the substrate and the incoming tows. The choice of heating method depends on the material being processed, with thermoplastic and thermoset materials requiring different approaches [23]. Thermosets typically require lower temperatures during layup, with Infrared (IR) heaters often being sufficient to provide the necessary tackiness without initiating premature curing. In contrast, thermoplastic materials demand much higher temperatures, often necessitating the use of laser heaters or other advanced systems capable of delivering precise and controlled heat. The goal is to reach the optimal interface temperature that facilitates intimate contact between layers while avoiding material degradation. Recent advances in heating technology, such as pulsed light systems, offer targeted and efficient heating solutions that enhance process control and part quality [24].

1.2.4.2 COMPACTION PRESSURE

Compaction pressure is crucial for achieving intimate contact between the layers of material in an AFP process. The compaction roller, typically featuring a metal interior and a silicone outer layer, applies pressure to ensure that the incoming tows adhere properly to the substrate and that voids between layers are minimized. While sufficient compaction pressure is essential for high-quality layup, excessive pressure can cause material degradation, particularly in thinner laminates [25]. For thicker laminates, the influence of compaction pressure diminishes, but it remains an important factor in the overall process. Innovations in compaction rollers, such as segmented and perforated designs, have improved the ability to conform to complex tool geometries and maintain consistent pressure across the layup.

1.2.4.3 SPEED

The speed at which material is laid down during AFP is another vital parameter. The challenge lies in balancing the need for high production rates with the requirement for high-quality layup [26, 27]. Lower speeds generally result in longer thermal exposure, which can improve polymer healing but may also increase the risk of material degradation if temperatures are too high [28]. Conversely, higher speeds reduce the dwell time, the time during which heat, and compaction pressure are applied, potentially leading to weaker cohesive forces and inferior interlaminar strength [29]. Therefore, speed must be carefully calibrated in relation to other parameters, such as heat and compaction pressure, to ensure a consistent and defect-free layup.

1.2.4.4 TOW TENSION

Tow tension is another parameter that plays a crucial role in the AFP process. Proper tension ensures that the tows are placed accurately and consistently, maintaining their shape and alignment as they are deposited onto the tool surface [30]. However, excessive tow tension can lead to issues such as tow slips and bridging, especially in concave areas of the tool [31]. Conversely, too little tension can result in inadequate placement, particularly when laying steered paths with in-plane curvature. The optimal tow tension varies depending on factors such as material properties and layup geometry, and achieving the correct balance is key to minimizing defects and ensuring the structural integrity of the final part.

1.2.4.5 ENVIRONMENTAL MONITORING

In addition to these primary process parameters, environmental factors such as temperature and humidity within the manufacturing environment must also be carefully controlled. Maintaining a stable environment is essential for consistent and high-quality manufacturing. For instance, ambient temperature fluctuations can affect the tackiness of thermoset materials and increase the risk of resin and fuzz deposits within the AFP head. Similarly, variations in humidity can impact material handling and adhesion. AFP processes are often conducted in cleanroom environments with stringent controls to minimize the presence of Foreign Object Debris (FOD) and ensure that the manufacturing conditions are optimal for producing defect-free parts.

1.2.5 EVOLUTION AND CURRENT ROLE OF AFP

The evolution of AFP technology has been a journey of innovation and adaptation spanning over seven decades. The inception of automation in composite production can be traced back to the 1950s and 1960s with the use of ATL and FW. Building on the tape-slitting concept introduced in Section 1.1, Hercules developed the first dedicated AFP machine around 1980 (Figure 1.8), with the first machine and manufactured parts produced between 1983–1984 [32]. These machines, commercially available by the end of the decade, presented a confluence of the differential payout capability of FW and the compaction/cut-restart capabilities of ATL. The AFP system's capability to modulate layup speed, pressure, temperature, and tow tension significantly enhanced its effectiveness, attracting aerospace OEMs such as Boeing, Lockheed, and Northrop as early adopters [33], and thus sparking the continuous innovation of the technology.

The years between 1985 and 2000 marked a period of significant advancements. Bullock's demonstration of an offline programming system enhanced machine production time, allowing independent programming that could later be uploaded to the machine for execution [34]. Offline programming provided flexibility compared to conventional programming, by allowing simulation testing of experimental and innovative programs prior to production. In 1993, a report by Grant and Benson outlined the implementation of a refrigerated creel system to resolve issues within the creel, extending material life and ensuring clean unspooling [35]. Toward the end of the century, AFP research focused on

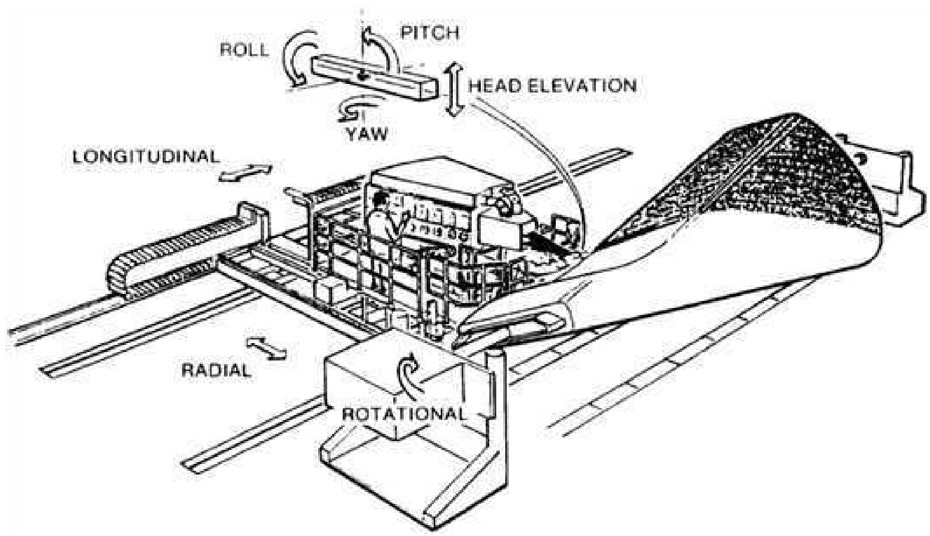

Figure 1.8: Drawing of early AFP machine created by Hercules [18].

productivity enhancements and the development of thermoplastic composites for aerospace structural applications. A system delivering up to 24 tows at once represented a turning point, as it reported a layup rate of up to 30 m min^{-1}, more than doubling the productivity associated with manual layup [36, 37].

The early 2000s saw a continued commitment to improving process reliability and productivity. Key areas of focus included tow steering [38, 39], thermoplastic materials [23, 40], and IR heating [24]. A 2006 patent by Engelbart et al. marked the inception of an automated inspection system [41], which Boeing demonstrated to reveal that inspection and rework accounted for 63 % of the cycle time required for manufacturing a fuselage barrel, more than 2.5 times the duration of the layup process itself [26]. Electroimpact found that inspection and repair consumed 32 % of the total time, while machine layup time was 27 % [27]. Electroimpact also contributed significantly during this time with a high-speed AFP system capable of 50.8 m min^{-1}, which featured interchangeable heads and a reduced tow-path length [28].

The period between 2010 and 2015 ushered in a new era of technological advancements. Notably, research by Elhajjar et al. in 2010 introduced efficient simultaneous use of multiple machine cells and modular AFP heads [29]. The modular head offered benefits such as 360-degree positioning, a variety of tow widths, short tow paths, and offline maintenance. Subsequent work by Jeffries et al. in 2013 demonstrated highly accurate robotic layup with a 3-sigma positional accuracy of ±0.08 mm [30].

Today, AFP is a vital tool in composite manufacturing, focusing on high throughput, minimal defect layups, and in-situ thermoplastic layups. In-situ thermoplastic layups eliminate the costly and size-limiting curing step, improving overall quality and efficiency [31]. Looking forward, digital twins of the AFP process, digital parts, and full lifecycle optimization represent key areas of future development [42]. The iSPP framework presented in this dissertation contributes to this vision by integrating design, manufacturing, and inspection into a closed-loop system. All the advancements presented in this section are summarized in Figure 1.9, demonstrating AFP's commitment to innovation and its persistent pursuit of efficient, quality manufacturing. Despite this progress, significant challenges remain; the following section examines the shortcomings that continue to limit AFP's broader adoption.

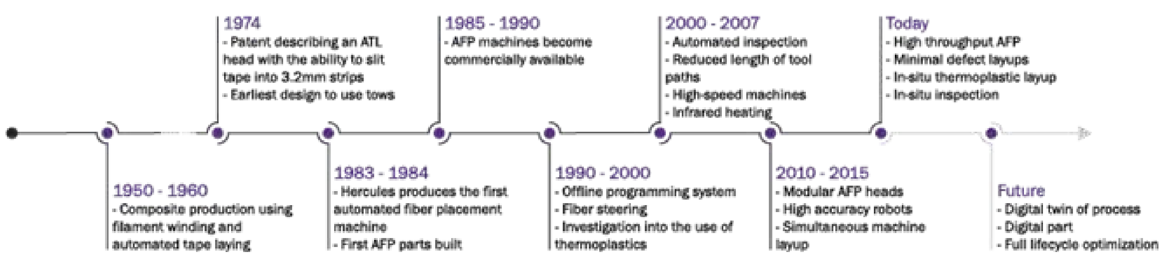

Figure 1.9: Timeline of AFP development.

1.3 SHORTCOMINGS OF AFP

Although AFP has become a cornerstone of advanced composite manufacturing, the process still faces significant limitations that restrict its broader adoption across industries. These shortcomings are primarily related to cost, process inefficiency, lifecycle integration, and inspection.

1.3.1 COST AND COMPLEXITY

AFP systems are capital-intensive, requiring specialized equipment, tooling, and facilities. The process also demands highly skilled engineers to perform process planning, operation, and quality assurance. As a result, AFP is often cost-prohibitive outside of aerospace, where the performance benefits can justify the expense. In industries such as automotive, maritime, and wind energy, the high cost of AFP remains a barrier to widespread adoption. Reducing this barrier requires tools that lower the expertise threshold and accelerate process planning decisions.

1.3.2 TRIAL-AND-ERROR PROCESS PLANNING

Current AFP process planning practices rely heavily on trial-and-error methods or legacy parameter sets [18]. Engineers often perform multiple trial layups to identify suitable process parameters for a new part, which increases material waste, cycle time, and cost. This empirical approach is inefficient and does not scale well as part complexity increases. Furthermore, parameter values that work for one geometry may not be transferable to another, limiting repeatability and consistency. The CAPP software developed in this dissertation addresses this limitation by semi-automating process planning decisions and providing quantitative scoring metrics that guide parameter selection.

1.3.3 SILOED LIFECYCLE STAGES

A major limitation of current AFP practice is the separation of design, process planning, manufacturing, and inspection into isolated stages [43]. In most industrial workflows, design decisions are made without full consideration of manufacturing constraints, while inspection data is rarely fed back into design or planning. This siloed approach prevents continuous improvement, reduces efficiency, and often results in suboptimal structures. Ideally, these stages should be interconnected, enabling feedback loops that improve manufacturability, reduce defects, and optimize structural performance. The iSPP framework presented in this dissertation directly addresses this gap by establishing bidirectional data exchange between design tools, manufacturing simulation, and inspection systems.

1.3.4 INSPECTION BOTTLENECKS

Inspection and rework represent one of the most significant bottlenecks in AFP. Studies have shown that inspection can consume more time than the layup process itself, with some reports indicating that inspection and repair account for more than 60 % of total cycle time [26]. Manual inspection is labor-intensive, time-consuming, and subject to variability based on operator expertise. While automated inspection technologies are emerging, they are not yet fully integrated into AFP workflows, and many systems still rely on human intervention to detect and repair defects. By incorporating inspection data into a closed-loop feedback system, the iSPP framework enables comparison of as-manufactured and as-designed conditions, supporting continuous process refinement.

1.3.5 LACK OF STANDARDIZATION AND INTEROPERABILITY

Another challenge is the lack of standardization across AFP machines, software platforms, and inspection systems. Each manufacturer often develops proprietary solutions, which creates barriers to interoperability and complicates integration into broader digital manufacturing ecosystems. This fragmentation slows the adoption of best practices and limits the scalability of AFP across industries. While this dissertation does not solve the standardization problem directly, the modular architecture of the iSPP framework demonstrates how existing tools can be integrated through defined interfaces, providing a template for interoperability.

In summary, AFP remains constrained by high cost, reliance on empirical process planning, siloed lifecycle stages, inspection bottlenecks, and a lack of standardization. These shortcomings limit its efficiency, repeatability, and accessibility. Addressing these challenges requires a shift toward smart AFP systems that integrate design, manufacturing, and inspection into a closed-loop, data-driven workflow. Such an approach has the potential to reduce costs, improve quality, and expand AFP's applicability beyond aerospace into industries such as automotive, maritime, and wind energy. The following section introduces the vision of Smart AFP and positions the iSPP framework as a comprehensive response to these challenges.

1.4 TOWARDS SMART AFP

The limitations of current AFP practice, outlined in the preceding section, highlight the need for a paradigm shift in how composite structures are designed, manufactured, and inspected. Rather than treating these stages as isolated processes, the next generation of AFP must integrate them into a unified, data-driven workflow. This vision, often described as Smart Automated Fiber Placement [11], builds on the principles of Industry 4.0 by combining digital models, real-time monitoring, and adaptive decision-making to create a closed-loop manufacturing system.

A smart AFP system is enabled by three interconnected infrastructures. The first is the physical infrastructure, which includes advanced sensors, compaction systems, and heating mechanisms that provide real-time data on material behavior and process conditions. The second is the cyber infrastructure, which connects machines, software, and operators through robust communication networks, enabling interoperability and digital twin integration. The third is the data infrastructure, which manages and analyzes large volumes of process and inspection data, supporting predictive models, defect detection, and continuous improvement. Together, these infrastructures establish the foundation for an intelligent AFP ecosystem.

By leveraging these capabilities, smart AFP systems can move beyond empirical trial-and-error methods and toward predictive, adaptive manufacturing. Several advancements are central to this transition. One is material-independent process characterization, where machines automatically determine optimal process parameters for different material systems without extensive manual testing. Another is automated process placement, in which toolpaths and layup strategies are optimized holistically to balance coverage, defect minimization, and efficiency. A third is integrated inspection and defect management, where in-situ monitoring and automated defect detection reduce reliance on manual inspection and enable real-time corrective actions. Finally, OoA processing offers the potential to reduce cost and expand part size beyond autoclave limitations by combining AFP with alternative consolidation and curing methods.

This vision aligns with the Department of Defense (DoD)'s emphasis on digital engineering, as outlined in DoD Instruction 5000.97 [44]. That policy defines digital engineering as the use of models, digital threads, and authoritative data sources to connect design, manufacturing, testing, and sustainment across the system lifecycle. Embedding these principles into AFP provides a clear framework for achieving interoperability, traceability, and continuous improvement, ensuring that manufacturing decisions are informed by accurate, model-based data.

This dissertation contributes to this vision by developing the iSPP framework for AFP. The framework unifies design, manufacturing, and inspection into a closed-loop workflow, enabling intentional control over defect placement, improved process planning efficiency, and the incorporation of structural performance considerations into manufacturing decisions. By embedding defect awareness, structural analysis, and inspection feedback into a single environment, the research presented here advances AFP toward a more intelligent, accessible, and cost-effective technology. The following section summarizes the specific contributions of this dissertation and outlines the structure of the remaining chapters.

1.5 DISSERTATION CONTRIBUTIONS AND OUTLINE

This dissertation advances AFP from a siloed, trial-and-error process into a smart, lifecycle-integrated system that unifies design, process planning, manufacturing, and inspection. The work aims to reduce cost, improve efficiency, and enable intentional control over defects and structural performance, thereby lowering the barrier to AFP adoption across industries beyond aerospace. The primary contribution is the development of the iSPP framework, which serves as the foundation for Smart AFP by enabling semi-automated process planning and integration with existing industrial tools. Building on this, the research introduces defect scoring and optimization methods to quantify manufacturing quality and guide decision-making at both the ply and laminate levels.

A key aspect of this work is the deep integration of design and manufacturing data, achieved by linking the CAPP software with structural analysis tools (HyperX), manufacturing simulation software (VCP), and in-situ inspection systems (RIPITx). To create a truly adaptive framework, the system incorporates closed-loop inspection feedback, allowing as-manufactured data to refine process planning and improve future builds. These contributions are demonstrated through application to the Quartic benchmark surface, which confirms the benefits of intentional defect placement, improved planning efficiency, and the potential for enhanced structural performance through margin-informed process optimization.

The specific contributions of this dissertation span software development, algorithmic formulation, and computational demonstration. The first contribution is the development of the CAPP software, a central integration platform that orchestrates data exchange between design, manufacturing, and inspection tools, providing the computational environment within which the subsequent algorithmic contributions operate. The second contribution is a defect scoring methodology comprising three novel metrics: the MIDS, which quantifies manufacturing quality based on defect frequency and severity; the Design-Informed Defect Score (DIDS), which incorporates structural margin-of-safety data from HyperX; and the IDS, which combines both perspectives through user-defined weighting to enable trade-off analysis between manufacturing efficiency and structural performance.

The third and fourth contributions address process optimization at different scales. PLO employs a two-phase algorithm combining global surveying of starting points and layup strategies with local refinement via perpendicular shift to identify optimal ply configurations. LLO extends optimization through the laminate thickness with two complementary methods: a stagger-based approach for systematic starting point distribution, and a gravity-based model derived from Newton's law of universal gravitation that quantifies defect interactions to minimize through-thickness stacking while protecting margin-critical regions. The fifth contribution is the incorporation of as-manufactured inspection data from RIPITx into CAPP, enabling comparison with as-designed predictions and supporting continuous process refinement through closed-loop feedback. Finally, the sixth contribution is the verification and demonstration of the complete iSPP framework through application to the Quartic benchmark surface, a doubly-curved geometry representative of challenging AFP manufacturing conditions. Each framework component is exercised in sequence confirming functional correctness and illustrating the integrated workflow from design intent through process planning.

The remainder of this dissertation is organized as follows:

Chapter 2 Literature Review — Reviews composite manufacturing methods, AFP technology, and prior work in process planning, defect management, and lifecycle integration, identifying the research gaps addressed in this dissertation.

Chapter 3 The iSPP Framework — Presents the iSPP methodology, including VCP integration for defect prediction, the MIDS, DIDS, and IDS scoring formulations, PLO and LLO algorithms, and inspection data incorporation. The chapter establishes CAPP as the central platform connecting design intent, process optimization, and quality verification.

Chapter 4 Case Study: The Quartic Benchmark Surface — Verifies and demonstrates the iSPP framework through systematic application to a doubly-curved benchmark geometry. The chapter presents component-level verification of the CAPP-VCP subprocess integration and HyperX data exchange, demonstration of MIDS, DIDS, and IDS scoring discriminability through a complete structural feedback loop, ply-level and laminate-level optimization results, and an inspection data visualization demonstration using representative RIPITx-format data.

Chapter 5 Conclusions and Future Work — Summarizes the contributions of this dissertation, discusses the implications for AFP process planning practice, and outlines directions for future research including enhanced material modeling, expanded optimization algorithms, and broader industrial application.

CHAPTER 2

LITERATURE REVIEW

The preceding chapter established that current AFP practice remains constrained by siloed lifecycle stages, trial-and-error process planning, and limited feedback between design, manufacturing, and inspection. Addressing these challenges requires a comprehensive understanding of the state of the art across each domain of the AFP lifecycle. This chapter provides that foundation through a systematic review of the literature, establishing the theoretical and practical basis for the iSPP framework developed in subsequent chapters.

The review is organized as follows. Section 2.1 establishes the broader lifecycle context, outlining the distinct yet interconnected phases of design, process planning, manufacturing, inspection, and post-processing (with post-processing serving as an auxiliary stage rather than a core lifecycle pillar). Section 2.2 examines the fundamentals of AFP process planning, detailing the critical decisions regarding layup strategy and process parameters that dictate manufacturing outcomes. Section 2.3 presents a detailed taxonomy of common AFP defects, examining their formation mechanisms, effects on structural performance, and modeling approaches. Section 2.4 surveys the optimization efforts developed to improve the process, spanning design optimization, path planning at the ply and laminate levels, process parameter tuning, and nascent efforts in Design for Manufacturing (DFM). Finally, Section 2.5 synthesizes the literature to identify the critical gaps in current practice that motivate the iSPP framework presented in this dissertation.

2.1 LIFECYCLE CONTEXT OF AFP

The AFP lifecycle encompasses a series of interconnected phases that transform a composite design into a finished structure. Understanding the roles, challenges, and interdependencies of these phases is essential for identifying opportunities to improve the overall process. This section provides an overview of each phase—design, process planning, manufacturing, inspection, and post-processing—and concludes with a discussion of the ideal closed-loop lifecycle that integrates these traditionally siloed activities.

2.1.1 DESIGN

The design phase is the foundational stage of the AFP lifecycle, establishing the final part's structure and performance characteristics. Decisions in this phase influence all subsequent stages of process planning, manufacturing, and inspection. This phase involves a comprehensive requirements analysis that defines the load conditions, environmental factors, and performance targets for the composite structure. Based on these inputs, engineers create the detailed design, specifying the geometrical shape, stacking sequence, ply boundaries, and fiber orientations.

In practice, composite designs often utilize conventional constant stiffness laminates, such as quasi-isotropic and balanced configurations with fiber angles restricted to 0°, ±45°, and 90°, which are guided by strict industry guidelines to ensure robustness. However, AFP technology also enables more advanced variable stiffness designs, where fiber paths are optimized to align with principal load paths, enhancing structural efficiency and performance. To validate these designs, engineers rely heavily on computational methods, particularly Finite Element Analysis (FEA), to simulate structural behavior and predict performance [45–47]. Recent advancements have also introduced Machine Learning (ML) and AI models to augment these simulations, accelerating ply optimization and predicting laminate behavior under complex loading scenarios [48]. This reliance on computational tools results in a relatively high Level of Automation (LoA) in the range of 5–6 for the design process itself [22].[1]

[1]The LoA scale, as defined by Brasington et al. [22], is a 1–8 ordinal measure of process automation.

A primary challenge in this phase is reconciling structural performance with manufacturability. While AFP allows for complex geometries and precise fiber steering, these designs must account for the physical limitations of the manufacturing process. Overlooking these constraints can lead to the formation of defects such as gaps, overlaps, wrinkles, and fiber angle deviations during production [49]. Currently, the design phase often operates in an open-loop or siloed manner, with limited systematic integration of downstream manufacturing realities or inspection feedback [43]. This disconnect necessitates manual, time-consuming iterations between design and process planning teams to mitigate manufacturability issues. Ideally, a closed-loop workflow would incorporate data and feedback from previous manufacturing cycles, allowing designers to proactively refine layouts, reduce defects, and improve overall part quality from the outset.

2.1.2 PROCESS PLANNING

The process planning phase in AFP serves as the critical bridge between the theoretical design and its physical manufacturing. This stage translates the finalized design specifications into a set of executable instructions for the AFP machine, accounting for material properties, equipment capabilities, and the geometry of the tool surface. The primary output of this phase is the generation of toolpaths that guide the machine's end effector.

At its core, process planning involves several key decisions, including the selection of layup strategies [7] and the choice of processing parameter windows [50]. Engineers typically select between reference curve strategies such as the fixed-angle (or rosette) method, which maintains a constant fiber angle, and the geodesic method, which follows the natural curvature of the surface to minimize steering. Once a reference curve is established, it is propagated across the surface using a coverage strategy, such as the parallel or shifted method. These choices are fundamental, as they directly influence the final fiber orientations and the formation of manufacturing defects. A detailed treatment of these process planning fundamentals is provided in Section 2.2.

Despite its importance, process planning is often a significant constraint in the AFP workflow. The process for selecting inputs remains a largely manual, trial-and-error endeavor that relies heavily on the experience and judgment of the process planner. This approach results in a low LoA of 2–3 and can lead to suboptimal and inefficient manufacturing plans [22]. While modeling and simulation tools exist to provide insight and allow for virtual validation [51], their use is often constrained by a trade-off between computational expense and predictive accuracy. Although research has explored the semi-automation of these tasks [52, 53], widespread industrial implementation is still limited. Consequently, process planning is frequently conducted as a discrete, siloed activity with insufficient feedback from manufacturing and inspection, highlighting a critical gap in achieving a fully optimized AFP lifecycle.

2.1.3 MANUFACTURING

The manufacturing phase is where the digital plans from design and process planning are physically realized. In this stage, a numerically controlled robotic or gantry-based system, equipped with a fiber placement end effector, executes the programmed toolpaths to deposit composite tows onto a tool surface or mandrel. This process is effectively a form of additive manufacturing, constructing the laminate ply by ply to form the final green state structure. The high degree of mechanical automation allows for precise material placement at speeds approaching $100\,\mathrm{m\,min^{-1}}$ on low-contour parts [54], resulting in a high LoA of 6–7 [22].

Execution of the manufacturing plan involves the simultaneous and carefully controlled application of several critical process parameters. The material, typically pre-impregnated tows stored in a creel, is fed through the machine head. As it is deposited, a heating system such as an IR lamp or laser raises the material temperature to achieve the necessary tack for thermosets or to facilitate consolidation for thermoplastics [55].

Concurrently, a compaction roller applies pressure to ensure intimate contact between layers, promote adhesion, and minimize voids [56]. The interplay of these parameters, along with tow tension and layup speed, is crucial for achieving the desired part quality.

Despite the high level of robotic precision, the manufacturing phase is often treated as a black box and faces significant challenges. The process is highly sensitive to variations in parameters, and the need for high throughput often creates a trade-off between speed and quality. Higher speeds reduce the material's thermal exposure and compaction dwell time, which can compromise layup quality and interlaminar strength. Furthermore, current digital twin implementations are often limited to kinematic simulations for collision avoidance rather than comprehensive, multi-physics process analyses [42]. There is a notable lack of real-time analysis or control loops to adjust for deviations. This operational isolation, without dynamic feedback from the actual manufacturing process or downstream inspection, underscores a major gap in achieving a truly intelligent and interconnected AFP system [11].

2.1.4 INSPECTION

The inspection phase is the fourth pillar of the AFP lifecycle, dedicated to quality control through the detection, characterization, and documentation of manufacturing defects. This stage is critical for verifying that the as-manufactured part meets design specifications and for mitigating the risks associated with a high-speed, automated process where defects can rapidly propagate across plies. These defects, which include positioning defects (gaps, overlaps), bonding defects (wrinkles, puckers), tow defects, and FOD, can have a deleterious effect on the final part's structural performance.

The current industry standard for inspection relies primarily on manual, visual assessment by trained operators [57]. However, this method faces significant challenges. The primary difficulty stems from the low visual contrast between the incoming tows and the substrate, making it difficult to reliably identify defects. Consequently, manual inspection is labor-intensive, time-consuming, and subject to operator variability, creating a major impediment to production efficiency. Studies have shown that inspection and rework can consume a substantial portion of the total cycle time, with estimates indicating that these activities account for as much as 63 % of the time required to manufacture a fuselage barrel while actual machine layup accounts for only 27 % [26, 58, 59]. Research has shown that manual rework itself can introduce new defects or redistribute existing ones rather than eliminating them, with one study reporting that gap area actually increased after rework and identifying a previously unreported defect type caused by the repair process [60].

To overcome these limitations, the field has moved toward automated Non-Destructive Inspection (NDI) techniques, many of which are suitable for in-situ monitoring [61]. Thermal imaging, or thermography, monitors the thermal profile of the laminate during layup to detect variations that indicate defects such as gaps, overlaps, and bonding issues [62–65]. Laser profiling, or profilometry, creates a detailed height map of the surface using laser projections to precisely identify and measure geometric defects such as wrinkles, folds, and puckers [59, 66–68]. Eddy current inspection is employed for conductive materials to detect subsurface inconsistencies and defects [69, 70]. These non-destructive methods offer improved speed and consistency, particularly in in-situ configurations that enable real-time monitoring with minimal disruption to the manufacturing process.

More recently, these sensor technologies have been coupled with ML and computer vision systems [71]. These advanced models can process image and sensor data in real time to automatically identify, classify, and even predict the severity of defects, often with accuracy and speed that surpass human capabilities. Despite these technological advancements, inspection largely remains a reactive and isolated process. The valuable data generated on defect occurrence, type, and location is not yet systematically fed back to influence and optimize upstream design and process planning decisions, reinforcing the open-loop nature of the current AFP workflow.

2.1.5 POST-PROCESSING

The post-processing phase is the final stage in the AFP lifecycle, where the green state laminate is consolidated and cured to achieve its final mechanical properties, dimensional accuracy, and surface quality. While not one of the four core pillars of the AFP process itself, this auxiliary phase is essential for completing the manufacturing workflow. The specific requirements of post-processing are dictated by the material system used—thermoset prepreg, thermoplastic prepreg, or dry fiber—with each presenting distinct process chains, challenges, and opportunities.

As the most common material in aerospace AFP, thermoset prepregs require a separate, high-cost curing step after layup. The standard procedure involves autoclave curing, where the part is subjected to elevated pressure (up to approximately 1 MPa) and temperature (typically 120 °C to 180 °C) for several hours. This process consolidates the laminate, minimizes voids, and fully cross-links the polymer matrix to achieve the desired strength and stiffness. While reliable, autoclave curing is a significant constraint in manufacturing. It is a costly, energy-intensive, and time-consuming batch process, and the physical size of available autoclaves places a hard limit on the scale of composite structures that can be produced [72].

Thermoplastics offer a compelling alternative by enabling in-situ consolidation [73]. Because thermoplastic resin can be repeatedly melted and solidified, consolidation can occur directly during the layup phase. Through precise, high-temperature heating (often around 400 °C) and immediate compaction at the nip point, the layers are fused together, potentially eliminating the need for a post-cure autoclave cycle. This approach offers a significant reduction in production time and cost. However, it requires a much narrower and more tightly controlled processing window to avoid material degradation, manage residual stresses from non-uniform cooling, and prevent part deformation. Some parts may still undergo a subsequent annealing step to relieve internal stresses and improve interlaminar bonding.

Structures made from dry fibers, which are not pre-impregnated with resin, require a multi-step post-processing sequence [74]. After the layup is complete, often facilitated by a thermoplastic veil to provide tack, the part must undergo a Resin Infusion (RIF) process such as RTM or VARTM. In this step, a pump is used to push liquid resin through the dry fiber preform, ensuring complete saturation. This is followed by a curing cycle in an oven to cure the matrix. A key advantage of dry fiber processing is that it enables OoA manufacturing, eliminating the need for expensive autoclave equipment and removing the associated size constraints on part geometry. While the lack of a tacky matrix can simplify steering during layup, the infusion process adds complexity to the manufacturing workflow, and realizing the full benefits of OoA processing requires careful management of this trade-off during both design and process planning.

2.1.6 FULL LIFECYCLE

The ideal AFP lifecycle is a seamless, interconnected workflow that integrates the four pillars of design, process planning, manufacturing, and inspection [43]. As introduced in Section 1.4, this vision of Smart AFP establishes a closed-loop system where data flows continuously between each stage, as illustrated in Figure 2.1 [11]. This communication fosters a cycle of continuous improvement, enabling the optimization of the entire manufacturing process.

The lifecycle begins with an initial, iterative loop connecting design and process planning. This collaborative phase aims to produce an optimal manufacturing plan by generating toolpaths that account for known manufacturing constraints from the outset. By simulating the process and anticipating potential defects such as gaps and overlaps, this step minimizes issues that would otherwise arise from isolated design decisions, ensuring that manufacturing realities are considered early.

The second phase, manufacturing, executes the optimized plan. During this stage, a substantial amount of data is collected from the machine, including process parameters (heating, speed, compaction) and sensor readings. This information feeds a high-fidelity digital twin or digital shadow, creating a detailed as-manufactured record of the part. This concept requires the integration of advanced, smart hardware capable of accommodating the complex and dynamic nature of the AFP process.

Following manufacturing, the inspection phase provides critical feedback by correlating the collected layup data with observed, real-world defects. This stage not only verifies product quality but also generates invaluable data that links specific manufacturing parameters to defect outcomes. This integration establishes a robust, data-driven foundation for understanding and predicting process behavior, which is essential for refining future manufacturing runs.

Finally, the inspection data closes the loop by feeding back into the design and process planning phases. This feedback enables the optimization of design parameters, such as ply boundaries, fiber paths, and material allowables, based on real-world manufacturing and inspection results. By completing the data pipeline, this stage establishes a continuous improvement cycle that transcends the conventional start-to-end workflow. Leveraging Industry 4.0 principles, this interconnected AFP workflow enhances efficiency, reduces defects, and unlocks new opportunities for ML and other advanced digital technologies, laying the foundation for truly intelligent and adaptive manufacturing systems.

2.2 AFP PROCESS PLANNING FUNDAMENTALS

Process planning is the critical link between the digital design of a composite structure and its physical realization through AFP. As introduced in Section 2.1.2, this phase encompasses two primary categories of decisions: layup strategy, which dictates the geometric placement of composite tows through the selection of starting points, reference

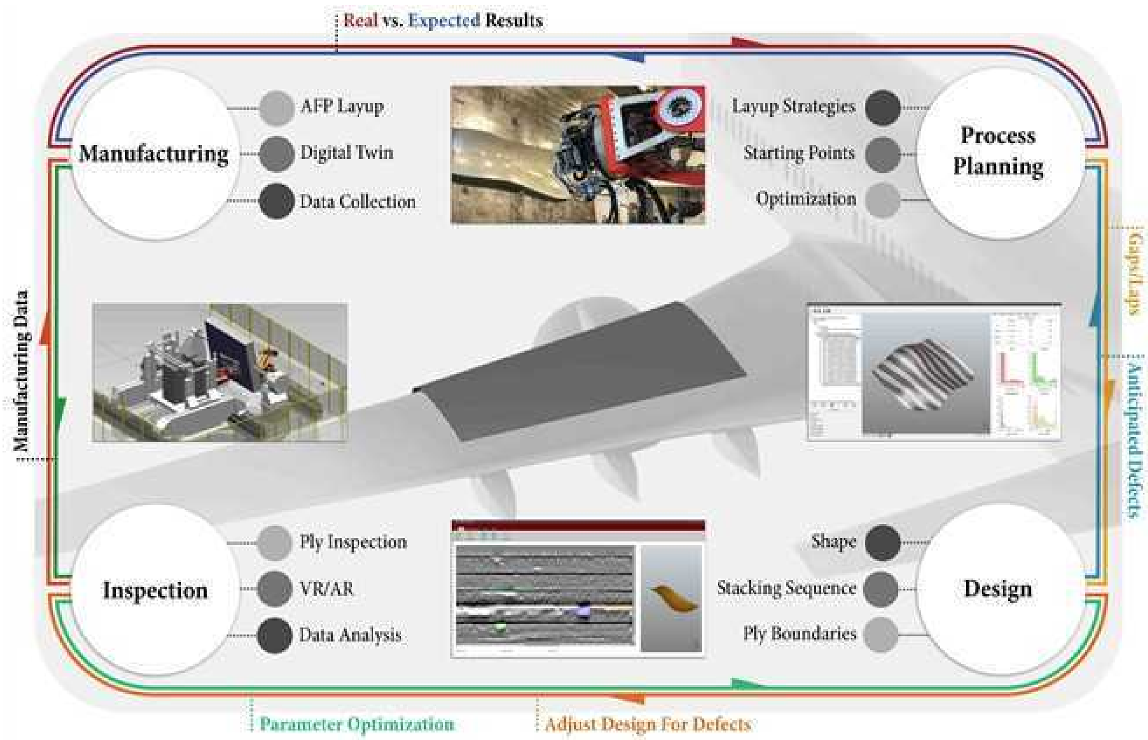

Figure 2.1: Graphical representation of closed-loop workflow [43].

curves, and coverage methods; and process parameters, which define the physical machine settings during deposition, such as roller pressure, heating temperature, deposition speed, and tow tension. The fundamental objective is to develop a manufacturing plan that adheres to the design's fiber orientation and structural requirements while minimizing defect formation and maximizing production efficiency. Critically, these two categories of decisions are interdependent and jointly govern the final quality and performance of the composite part, making a detailed understanding of their interplay essential for effective process planning.

2.2.1 LAYUP STRATEGY

Layup strategy in AFP defines the geometric approach for generating the toolpaths required to cover a specified ply boundary on the tool surface. It translates the design's intended fiber orientations into a set of toolpaths for the machine to follow. This strategy is governed by three primary decisions: the selection of a starting point, the definition of reference curves, and the application of a coverage strategy. The starting point establishes the initial tow deposition location, reference curves serve as the primary guides that establish the nominal fiber orientation across the ply, and coverage strategies propagate subsequent toolpaths from these curves to fill the entire ply boundary. The interplay of these choices directly dictates the distribution of geometric defects, such as gaps, overlaps, steering, and fiber angle deviation, particularly on complex, doubly curved surfaces. An effective layup strategy is therefore fundamental to achieving both manufacturing feasibility and structural performance, representing a key area of focus in process planning optimization.

2.2.1.1 STARTING POINT

The selection of a starting point represents the initial decision in toolpath generation, establishing the origin from which all subsequent paths are defined. This point is used to generate a reference curve, where a toolpath propagates from this origin based on a defined strategy. When a guide curve is provided, the starting point dictates its initial placement, with the final path for the first course being computed by interpolating that curve onto the tool surface from the specified point. Regardless of the method, the placement of this initial point directly influences the geometry of the resulting paths and, consequently, the location and magnitude of defects across the laminate. Depending on the selected point, the resulting toolpaths can exhibit significant variations, leading to unintended defect formation when comparing the as-manufactured part to the as-designed ideal.

The determination of the starting point's optimal location is a non-trivial task [75]. With no standardized methodology available, process planners in industry typically rely on a manual, trial-and-error process. This approach involves iteratively selecting a point and visually evaluating the performance of the resulting ply, often concluding when a solution is deemed qualitatively acceptable rather than proven to be optimal. Careful consideration during this process is required to ensure that the selection does not generate unnecessary defects or position them in structurally critical regions. Furthermore, manipulating the starting point location, often by shifting or staggering it between plies, is the primary mechanism for preventing the coalescence, or stacking, of defects through the laminate thickness [76, 77]. This makes the selection of the starting point a critical process input for controlling part quality and manufacturing feasibility.

2.2.1.2 REFERENCE CURVES

Once a starting point is established, a reference curve is generated to define the initial fiber path and constitutes the primary guide for toolpath generation. Its geometric properties dictate the nominal fiber orientation and the distribution of defects across the surface. The selection of a reference curve strategy is a critical process planning decision that involves a trade-off between maintaining the designed fiber angle and minimizing manufacturing defects. The literature identifies three principal types of reference curves: fixed angle, geodesic, and variable angle [7, 19, 78].

Fixed angle reference curves, also known as rosette or constant angle paths, follow a constant orientation relative to a chosen direction on the tool surface. This approach is widely used in industry because it provides predictable control over fiber angles, which is advantageous for flat or mildly curved parts. However, on highly complex or doubly curved geometries, forcing a constant angle path can deviate from the surface's natural geodesic path, necessitating significant fiber steering and leading to defects like wrinkling, gaps, and overlaps [79].

Geodesic reference curves, or natural paths, represent the shortest distance between two points on a 3D surface [80]. By following the natural contour of the tool, geodesic paths require minimal or no steering, thereby reducing the formation of steering-induced defects. This strategy is often preferred for complex surfaces where fiber alignment is difficult to maintain. The primary drawback is that the fiber angle may vary along the path relative to the design's coordinate system, which can result in deviations from the intended mechanical performance [81].

Variable angle reference curves allow for continuous changes in fiber orientation along the path, enabling the design of variable stiffness panels [82]. This method offers designers greater control over the mechanical properties of the laminate by tailoring stiffness and strength to specific load paths [83, 84]. While this approach provides significant performance benefits, it requires sophisticated path-planning algorithms, advanced computational resources, and precise manufacturing control to avoid defects, making it the most complex of the three strategies [85].

2.2.1.3 COVERAGE STRATEGIES

Once a reference curve is defined, a coverage strategy is employed to propagate toolpaths across the entire ply boundary, ensuring complete material coverage. This strategy dictates how subsequent courses are generated relative to the initial reference curve or each other. The choice of coverage strategy directly influences manufacturing outcomes, including the uniformity of spacing, the formation of geometric defects, and overall production efficiency. The literature details three primary strategies for surface coverage: offset curves, shifted curves, and independent curves [7].

The offset, or parallel, curve strategy is the most widely used method in process planning. In this approach, adjacent toolpaths are computed to be equidistant from a reference curve, which minimizes the formation of gaps and overlaps between courses. This can be accomplished through two main techniques: operating directly on the surface's parametric definition or on a discretized mesh. Parametric methods compute offsets numerically [86–88], as closed-form solutions are generally not available for 3D surfaces. One common technique generates an offset by stepping along the reference curve, moving along a vector normal to the curve on the surface's tangent plane, and then projecting the new point back onto the mold surface [86]. While computationally direct, this method introduces a geometric error that increases with both surface curvature and offset distance. To improve precision, a more accurate parametric approach involves intersecting the mold surface with a plane perpendicular to the reference curve and finding the new point by measuring the required geodesic distance along this intersection arc [86, 87, 89, 90]. Alternatively, mesh-based approaches, such as the Fast Marching Method (FMM) [91, 92], solve the problem on a discretized surface. FMM treats the reference curve as a wavefront at time $t = 0$ and propagates it across the mesh nodes based on the principles of the Eikonal equation. The final offset curves are then generated as the iso-contours of the calculated time values. While offset strategies excel at maintaining uniform spacing, their application to highly complex surfaces can cause the fiber orientation in subsequent paths to deviate significantly from that of the initial reference curve.

The shifted curve strategy involves simply translating the reference curve along its perpendicular direction on the surface [81, 93–102]. This method is valued for its simplicity, as each generated toolpath has the same geometry as the reference curve. However, on complex surfaces, this simple translation does not guarantee that the paths will remain equidistant, which often leads to the formation of gaps and overlaps where the surface curvature changes [97].

The independent curve strategy involves generating each toolpath without influence from the preceding ones. This approach is commonly applied to highly complex tool surfaces, where courses can be arranged in a staggered manner, maintaining constant length and varying directions [86]. The flexibility of generating each curve independently can be advantageous, but because the courses are not necessarily parallel, this method can induce significant gaps and overlaps. Furthermore, strategies based on this principle can create many short, discontinuous courses, which may reduce the resulting laminate's structural strength [103].

2.2.1.4 ADDITIONAL CONSIDERATIONS

Beyond the primary layup strategy, process planning involves defining a range of ply- and course-level parameters to ensure the design is manufacturable and meets performance requirements [52]. At the ply level, planners assess and finalize definitions for ply boundaries, internal cut-outs, and ply angles. Additionally, the inclusion of "dog ears," or small tabs of excess material at ply corners, is often evaluated. This feature is a manufacturing necessity for off-axis plies (e.g., ±45°) to accommodate the AFP machine's minimum tow length, which is a physical constraint based on the distance between the tow-cutting mechanism and the deposition point. Finalizing these geometric definitions often requires multiple iterations between designers and process planners, as any change necessitates further structural analysis to re-evaluate the part's performance.

Within each ply, planners also specify tow- and course-level settings. While the material type is dictated by the design, tow width can be adjusted based on surface geometry; more complex surfaces typically require narrower tows for better conformability, whereas simpler geometries can utilize wider tows or tapes to improve deposition rates. Planners also define the course width, which may vary across a ply to ensure adequate compaction, and evaluate course efficiency.

2.2.2 PROCESS PARAMETERS

Process parameters play a fundamental role in defining the quality and efficiency of AFP, necessitating careful selection and balance to achieve desired outcomes that may be conflicting, such as the trade-off between speed and quality. The proper selection and control of these parameters profoundly influences the final quality of the layup. Additionally, they carry significant implications for the resultant mechanical properties of the composite part, thus emphasizing their crucial role in the overarching AFP process [83, 104]. With industry's demand for high layup speeds, the calibration of these parameters often requires balancing layup quality against throughput [105].

The following subsections provide an overview of the primary processing parameters intrinsic to AFP. Each parameter is discussed in terms of its specific role, influence, and the effects that variations in its settings may impart on the process outcome as well as on other parameters, as they interact with each other. By doing so, the complex interplay of these parameters and their impact on the final product are illustrated.

2.2.2.1 PRESSURE

Compaction pressure is an instrumental parameter in the AFP process, as it impacts the final quality of the manufactured structures [106]. Pressure is applied via a rubber compaction roller to ensure proper adherence between the incoming tows and the substrate and to remove voids in the laminate [107]. However, while optimal pressure plays an essential role in fostering intimate contact between plies, over-compaction can lead to material degradation [108]. In thicker laminates, the influence of compaction pressure noticeably diminishes [109]. This parameter does not operate in isolation; it shares a complex relationship with several factors. These include the roller's inherent characteristics, applied temperature, layup speed, and the material's tack level, which together shape the compaction process [25, 49, 56, 110]. Mismanagement of these factors can result in defects such as gaps, overlaps, wrinkles, and bridging. Not only are these defects detrimental to the part's mechanical properties, such as strength and stiffness, but they also disrupt the manufacturing process's efficiency due to the need for manual repairs.

Given the critical role of compaction pressure in determining AFP process output, in-depth research has been conducted, producing insights essential for process improvement. One study evaluated the performance of five distinct compaction rollers, emphasizing their stiffness and structural design [56]. By employing both static tests and FEA, this research explored roller deformation and the pressure distribution across its width. The study's findings offered key insights into the effect of roller architecture on the resultant steered courses and associated defects. Another investigation centered on the nuances of roller deformation on non-uniform surfaces, particularly focusing on chamfered sandwich structures [111]. Using pressure-sensitive films, the study compared measured compaction with FEA simulations, revealing roller deformation dynamics across various surfaces in static setups. A theoretical model addressed compaction pressure on irregularly contoured surfaces [112]. Validating the model with empirical measurements, compaction was examined at four distinct locations on a winglet mold. Comprehensive research examined the intricacies of process parameters' influence on compaction during layup on curved surfaces in AFP [113]. Though compaction pressure might appear straightforward, it is deeply interwoven with numerous factors, making its management in AFP a challenging task. These studies are instrumental in demonstrating the complexities associated with compaction pressure in AFP, providing refined understanding of the process and helping to improve process efficiency and final part quality.

2.2.2.2 TEMPERATURE

Temperature plays a crucial role in the development of interlaminar strength by ensuring an optimal interface between the incoming tows and the substrate [114]. The layup temperature is a predominant factor in achieving the right degree of tack, which is pivotal for minimizing manufacturing defects [115–117]. While reaching an adequate temperature is crucial, it is equally important to avoid excessively high temperatures, as they can lead to material degradation. The degree of tack is instrumental in the formation of most layup defects within thermoset materials, with the layup temperature being its primary influencer [115–117]. Achieving a balance between temperature and tackiness is paramount, as a higher degree of tack is beneficial for securing the prepreg to the tool surface and ensuring its adhesion to subsequent plies [118]. Determining the ideal temperatures usually involves a resource-intensive process of trial and error [119, 120]. Nevertheless, tack characterization can offer a valuable starting point for setting the right applied temperatures [119].

For thermoplastic materials, the process involves heating the material beyond its melting point. This is then followed by consolidation through applied pressure and eventually solidification as the material cools [121, 122]. It is imperative to closely monitor the processing temperature to prevent it from becoming excessively high, as this can compromise the material's integrity, leading to degradation [123, 124]. Additionally, the nuances of temperature regulation pose challenges to the part's quality. Uneven temperatures across parts can lead to non-uniform cooling rates, resulting in residual stresses and subsequent part deformation [125]. Notably, adjusting the temperature also impacts elements like void dynamics [124, 126, 127], the nuances of material healing [128], and the achievement of intimate contact [129].

2.2.2.3 SPEED

A primary advantage of AFP is the increased production rates and manufacturing cost reductions compared to traditional hand layup techniques [130]. However, the aspect of speed in the AFP process has multifaceted implications. From the quality of the final composite part, through the thermal exposure of the material, to the overall productivity of the process, the speed of layup plays a central role. Despite limited literature in the domain of speed as it relates to AFP, studies have been performed examining each of these implications. Various layup velocities show alterations in layup quality and required processing parameters [131, 132]. Regarding thermal exposure, slower speeds contribute to extended heat exposure, fostering improved polymer healing until the point where the applied temperature begins to degrade the material [133]. Conversely, increasing the layup speed curtails the duration for which compaction force and temperature interact with the material, potentially compromising the material's cohesive forces [74].

2.2.2.4 TOW TENSION

The landscape of literature examining tow tension in the context of AFP is notably sparse, despite tow tension being a vital parameter in the process. Proper tension ensures accurate fiber placement and reduction of both positioning and bonding defects [134]. However, if the tow tension exceeds optimal levels, it can have adverse effects, primarily leading to the slipping of tows [135]. This slippage is a direct consequence of the tension force surpassing the adhesive force that maintains the tows in place. Given the limited research in this area, establishing optimal tension settings remains largely dependent on empirical testing and operator experience for specific material and geometry combinations.

The layup strategy and process parameter decisions detailed in this section directly influence the formation of manufacturing defects. The following section provides a detailed taxonomy of these defects, examining their formation mechanisms, effects on structural performance, and modeling approaches.

2.3 AFP DEFECTS: TAXONOMY, FORMATION, EFFECTS, AND MODELING

AFP permits faster and more precise deposition of material compared to hand layup. However, this increased speed, in conjunction with inadequate process planning, can lead to rapid defect production. These defects are often unwanted and can have a negative influence on the structure's performance [136, 137]. AFP defects can be categorized into four main groupings: (1) positioning defects, (2) bonding defects, (3) tow defects, or (4) foreign bodies [62]. Each has unique characteristics in significance, creation mechanism, and effects. Harik et al. provide a viewpoint modeling investigation of AFP defects [49]. Categorizing the defects in this way allows them to be defined from the viewpoint of the designer, the process planner, the machine operator, and the inspector. This thorough categorization of AFP defects provides an understanding of the importance of defects from four different perspectives: the cause, anticipation, existence, and significance. A comprehensive list of all defect types and their categories is given in Table 2.1. Brasington et al. expand on this investigation with the development of printable 3D models to provide an educational tool for visualizing each defect [138]. This exposes personnel to defect identification, significance, and effect, traditionally only available through hands-on experience with inspecting numerous AFP manufactured plies.

Due to the deleterious effects on the overall structural performance of the part [139–141], manual inspection and rework are required on a ply-by-ply basis. The rework process has been examined in depth in several studies with respect to the effect on cycle time and overall machine utilization [27, 59, 142]. Sacco et al. uniquely examine the results of rework on a complex AFP part to determine where rework helped, hindered, or was indifferent to the overall process [60]. The data is collected and analyzed using an automated inspection software developed previously at the University of South Carolina's McNair Center [67, 68, 71, 143]. In addition, a new class of defect specifically associated with the manual rework process was presented, giving a clear example of where rework can degrade quality.

While the taxonomy of AFP defects is extensive, this review concentrates on a specific subset of geometrically predictable defects that are central to process planning optimization. These include gaps, overlaps, fiber angle deviation, and steering-induced defects. This specific group is prioritized because their formation is fundamentally linked to the layup strategy and toolpath geometry, making them predictable and, therefore, optimizable during the planning phase. Furthermore, the through-thickness interaction of these defects, or defect stacking, is examined as it represents a critical failure mechanism that LLO seeks to mitigate. The following subsections provide a detailed analysis of the formation, effects, and significance of each of these defect types.

Table 2.1: List of defect types and their associated category [49].

Defect	**Category**	**Causes**	**Significance**
Gap/overlap	1	Fiber steering, layup over complex surfaces	Site for failure initiation, resin/fiber rich areas, site for wrinkling
Twist	1	Initiated by folding, rotation during bi-directional layups	Increase/decrease in local thickness
Missing tow	1	Discontinued material feeding, insufficient tack adhesion	Local thickness variations, resin rich pockets
Boundary coverage	1	Material cannot perfectly meet at edge of part	Affects shape of part, failure points if not trimmed
Angle deviation	1	Incorrect roller coverage, small steering radii	Causes overlaps, leads to resin rich areas
Wandering tow	1	Unsupported portions of tow between roller and cutter	Leads to gaps and overlaps
Position error	1	Obstruction of tow during feeding, incorrect machine reference, machine control issues	Results in gap, site for failure initiation, more pronounced influence since close to boundary
Fold	2	Tensioner errors, long or complex tow paths, steered paths	Substantial influence on local fiber volume fraction, creates resin rich areas
Pucker	2	Excess tow feeding	Significant loss of strength
Wrinkle	2	Tow placement at small steering radii	Causes gaps and folded tows, loss of strength
Bridging	2	Too much tow tension, insufficient tack adhesion	Resin rich areas, delamination
Loose tow	2	Length of tow is shorter than length between roller and cutters	Results in gaps/overlaps and missing tows
Splice	3	Two tows joined end to end during the slitting process	Local thickness change, site for failure initiation especially under compressive loads
FOD	4	Resin or fiber fuzz collects on head, other debris from production area	Improper adherence of next ply

2.3.1 GAP AND OVERLAP

Gaps and overlaps are the most commonly occurring defects in AFP and are a direct consequence of toolpath geometry interacting with the tool surface. A gap is formed when adjacent courses diverge, leaving a resin-rich area without fiber reinforcement between them. Conversely, an overlap occurs when adjacent courses converge, causing one tow or course to be placed on top of another, resulting in a localized fiber-rich region with a buildup of material. A 3D graphical representation of these defects is provided in Figure 2.2. The primary cause of these positioning defects is the inherent geometric challenge of maintaining a constant course width while steering fibers over complex, doubly curved surfaces, particularly when using a parallel coverage strategy. This makes their occurrence predictable during the process planning phase, as they are a direct result of the selected layup strategy.

The formation of these defects introduces local variations in laminate thickness and induces out-of-plane fiber waviness in adjacent plies [144]. During consolidation, material from subsequent plies is forced to flow into the void of a gap or drape over the excess material of an overlap. This fiber deviation is a primary driver of mechanical performance degradation. Studies have shown that overlaps can induce localized fiber waviness that diminishes load-bearing capacity, while gaps lead to resin-rich pockets that become sites for crack initiation and delamination [145, 146].

The structural impact of gaps and overlaps has been extensively studied. Experimental investigations have demonstrated that these defects can reduce compressive strength by anywhere from 5 % to 27 %, depending on their size, location, and the laminate's stacking sequence [147]. Research by Croft et al. further highlighted that the effects are more pronounced at the laminate level, with strength reductions up to 15 %, compared to less than 5 % at the individual lamina level [136]. This body of literature confirms that gaps and overlaps are not merely cosmetic flaws but are significant structural features that must be managed during process planning to ensure the integrity of the final part.

Figure 2.2: 3D graphical representation of a gap/overlap [18].

2.3.2 ANGLE DEVIATION

Angle deviation is a geometric positioning defect defined as the difference between the as-manufactured fiber orientation and the as-designed, nominal angle for a given ply, as illustrated in Figure 2.3. Unlike gaps and overlaps, which are voids or excesses of material, angle deviation represents a violation of the design intent, which can compromise the engineered mechanical properties of the laminate. Its formation is directly linked to the toolpath generation process, making it a predictable defect that can be quantified and optimized during process planning [148, 149].

The primary causes of angle deviation stem from the challenges of laying a flat tow onto a complex surface. It can be caused by incorrect roller coverage, where the compaction roller fails to maintain full contact, or by steering at a small radius, which can cause the tow to shift slightly from its intended path after deposition [150]. In some cases, angle deviation is a necessary concession made by the layup strategy itself to generate manufacturable toolpaths across a surface. Research has also highlighted a fundamental trade-off in process planning: certain layup strategies, such as a rosette parallel approach, may be selected to reduce the total area of gaps and overlaps at the known cost of increasing angle deviation [7].

Because angle deviation is a comparison between the actual and intended states, it cannot be identified by visual inspection alone and requires post-processing of the toolpath data against a design benchmark [71, 151]. The effects of this defect are significant, as deviations from the specified fiber angles can lead to unintended structural responses. Furthermore, angle deviation can cause localized overlaps where a deviated course is covered by a subsequent one, leading to the formation of resin-rich areas and potential sites for delamination.

Figure 2.3: 3D graphical representation of angle deviation [18].

2.3.3 STEERING

In AFP, steering is the intentional curving of the toolpath to follow complex surface geometries while attempting to maintain the desired fiber orientation. It is a fundamental process planning input rather than an inherent defect. However, the physical limitations of steering a finite-width tow are a primary source of several predictable manufacturing defects, most notably wrinkles and puckers [88, 152–154]. A steering defect is therefore typically quantified as a violation of a predefined steering radius constraint, which is set to prevent the formation of these associated flaws.

The formation of steering-induced defects is a direct result of the path length differential across the width of the tow. When a course is steered, its inner edge travels a shorter path than its outer edge. This mismatch causes the inner edge to experience compression, which can lead to out-of-plane buckling in the form of puckers (a local lift-off of the tow, shown in Figure 2.4a) [155, 156] or wrinkles (a wavy pattern of puckering along the tow's edge, shown in Figure 2.4b) [157, 158]. These out-of-plane defects can significantly reduce the strength of the laminate if not properly compacted. The decision to steer a path, therefore, involves a critical trade-off between adhering to the nominal fiber angle and avoiding the formation of these geometric instabilities.

To manage this trade-off, process planners define a minimum steering radius as a key constraint. This radius represents the tightest curve a tow can follow without buckling, and it is dependent on the material properties, tow width, and machine capabilities. Process planning software can then predict steering violations by comparing the curvature of the generated toolpath against this minimum allowable radius. By treating steering as a predictable and optimizable parameter, planners can adjust layup strategies to minimize the occurrence of wrinkles and other related defects, ensuring the manufacturability of the part [153].

2.3.4 DEFECT STACKING AND INTERACTION

Defect stacking, or through-thickness defect interaction, occurs when defects in multiple plies coalesce or overlap in the same geometric location within a laminate. While an isolated defect on a single ply can introduce a localized anomaly, the stacking of defects through the thickness can amplify their negative effects, creating a much more significant structural flaw [140]. For example, as illustrated in the Finite Element (FE) representation in Figure 2.5, stacked gaps can form a large, continuous resin-rich pocket, while stacked overlaps can create a severe local thickness buildup. These compounded features drastically exacerbate out-of-plane fiber waviness and are a primary driver of significant reductions in mechanical performance [147, 159].

Literature has consistently demonstrated that the structural impact of stacked defects is more severe than that of isolated ones. Studies have found that laminates with stacked

(a)

(b)

Figure 2.4: 3D graphical representation of a (a) pucker and (b) wrinkle [18].

gaps and overlaps exhibit significantly greater strength knockdowns than those with defects distributed across a single ply, with compression strength reductions reaching as high as 25 % in some cases [160]. This highlights the complexity of defect interactions, where the overall effect is dependent on the type of defects being stacked, their proximity, the number of pristine plies between them, and the overall stacking sequence of the laminate.

To mitigate the detrimental effects of stacking, the primary strategy employed in process planning is staggering [94, 161]. Staggering is the practice of intentionally offsetting toolpaths between plies, typically those with a common fiber angle, to ensure that defects do not align vertically. By dispersing defects throughout the laminate, staggering minimizes the severity of localized thickness variations and reduces out-of-plane fiber waviness. The challenge lies in finding an optimal staggering strategy across the entire laminate, considering interactions between all plies. This makes the strategic management and minimization of defect stacking a principal objective of LLO [76, 77].

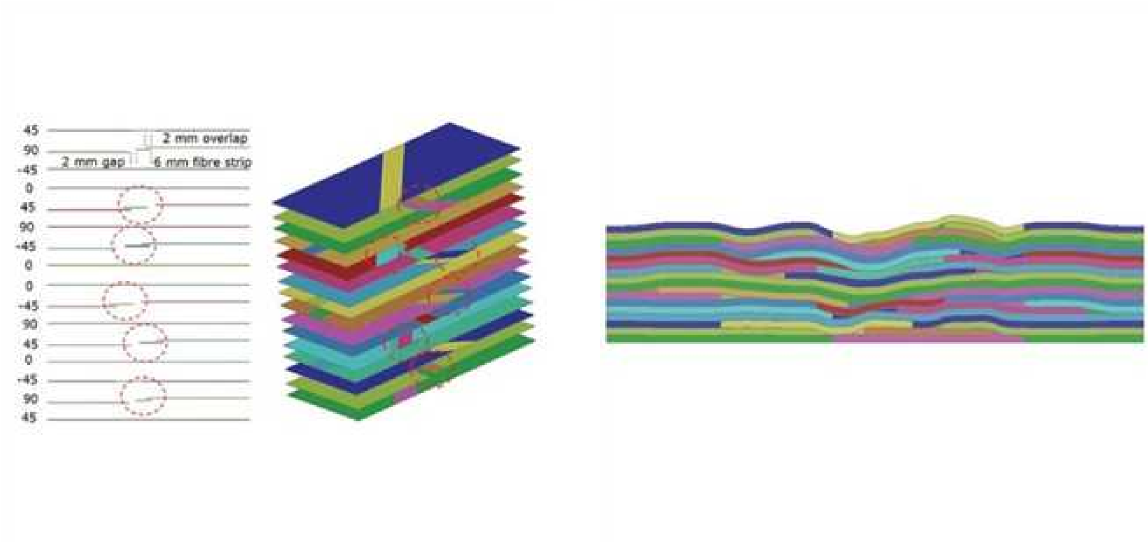

(a) (b)

Figure 2.5: Defect stacking FEM representation of (a) defect distribution and (b) cut-section from [159].

2.3.5 DEFECT MODELING

Due to the inherent complexity of the AFP process, defect occurrence is inevitable throughout the layup. These manufacturing defects can have a significant negative influence on a structure's performance; thus, it is vital to understand the creation and effect of each defect. The main defect types that are modeled are gaps, overlaps, angle deviation, and wrinkling resulting from steering. Essentially all other defect types can be represented by some modification or combination of these four types.

2.3.5.1 GEOMETRIC MODELING OF GAPS AND OVERLAPS

Gaps and overlaps are a result of the convergence or divergence of sequential fiber or tow paths. Overlaps are found as the intersection of any combinations of tow surfaces with the total area represented by:

$$A_{\text{overlap}} = \sum_{i,j} \text{tow}_i \cap \text{tow}_j \tag{2.1}$$

where tow_i and tow_j represent any combination of intersecting tows in the layup. Gap areas are represented by the total area of the tool surface minus the union of all tow surfaces:

$$A_{\text{gap}} = A_{\text{surface}} - \sum_{i,j} \text{tow}_i \cup \text{tow}_j \tag{2.2}$$

This results in any portion of the surface not covered by material. Both defects result in local thickness variations, with gaps causing resin-rich regions and overlaps causing fiber-rich regions, as shown in Figure 2.6. The limit to the possible size of a gap or overlap can also be established. The maximum allowable gap or overlap is usually a single tow width, after which a tow will be added or dropped to alleviate the defect. Gaps and overlaps are extracted geometrically, which can limit the predictive capabilities since other factors such as tow deformation are not introduced. Modeling of these defects is accomplished through FEA models and coupon testing, with the desired results being the change in laminate properties.

2.3.5.2 EXPERIMENTAL STUDIES ON MECHANICAL EFFECTS

Croft et al. performed a study in which the effects on ultimate strength due to gaps and overlaps were assessed through a series of tests [136]. These tests included tension, compression, in-plane shear, open hole tension, and open hole compression. The results

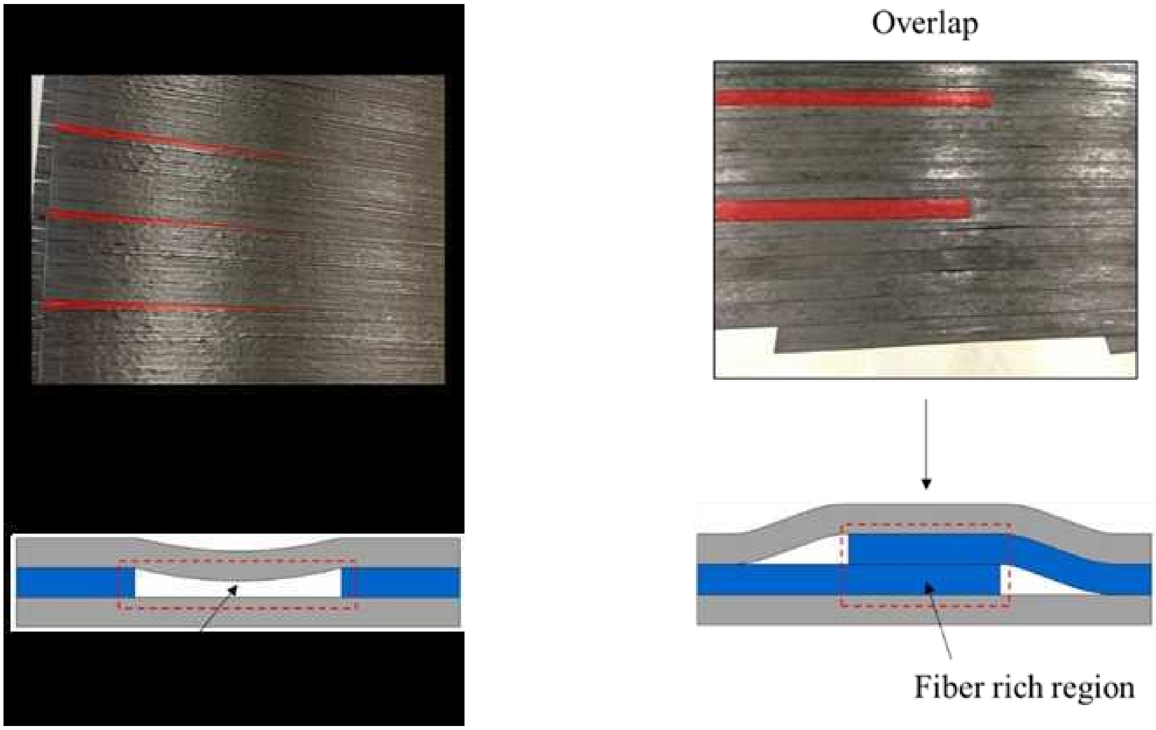

(a) (b)

Figure 2.6: Example geometries of (a) gap and (b) overlap defects [22].

are summarized in Table 2.2. The defects were placed in the middle layers of a symmetric laminate with the defects defined as one tow width and two tows thick. The reasoning for using a two-tow thickness was to increase the ability to capture the variability between defect configurations. This highlights the increasing effects when stacking defects through a laminate; however, this does not represent the effect of a defect in a single ply.

Nguyen et al. also investigated the effects of gaps and overlaps, termed the AFP manufacturing signature, with a defect size set of {12.7, 6.35, 3.18, 1.59, 0.794} mm widths [160]. The largest size represents a missing 12.7 mm tape while the smallest is a half-tow gap. It was found that defects below 6.35 mm had marginal effects on the strength properties of the laminate. Defects greater than 6.35 mm, which would be at least a tow width or more, showed more significant effects. Anay et al. also found that defects with small

Table 2.2: Summary of the effect of different defect configurations on ultimate laminate strength (reproduced from [136]).

		Gap	Overlap	Half Gap/Overlap
Tension		=	=	−
Compression		=	+	=
In-Plane Shear	Length	=	=	−
	Width	+	−	−
OHT		=	=	=
OHC	Length	+	+	+
	Width	−	−	−

"+": ≥ 3 % increase (up to 13%), "=": ±3 % variation, "–": ≥ 3 % decrease (up to 12%)

geometry have little to no correlation with the resulting strength and stiffness properties [162]. Marginal decreases in stiffness were found with gap defects and vice versa for overlaps. Additionally, overlaps cause a slight increase in ultimate strength, but gaps showed minimal effects.

Del Rossi et al. studied the effect of half gap and half overlap, meaning a gap and overlap with a half-tow width (3.175 mm in this case), on the strength of the resulting structure [163]. The length of the defect was 63.5 mm and the widths varied between −2.54 mm to 1.27 mm. The geometry of this defect type is a closer resemblance to that typically seen during AFP manufacturing. The defects were also distributed throughout the laminate at varying fiber angles to emulate AFP defects in multiple plies. For tensile specimens, strength reductions were usually smaller than 5 %, while compression specimens saw reductions of up to 20 % for the largest defects. The authors also noted that defects tend to have more effect when loaded in the matrix direction. It was also shown that the healing process of closing gaps and consolidating overlaps with caul plates is dependent on the stacking sequence and location of defects and does not necessarily further decrease the maximum strength.

2.3.5.3 FINITE ELEMENT ANALYSIS OF DEFECTS

Finite element analyses can also be employed to investigate the effects of gaps and overlaps. It is important to note that this type of analysis requires verification with experimental data to establish the efficacy of the models. Blom et al. developed a model to study the influence of tow drop areas, which result in gaps and overlaps, on stiffness and strength of variable stiffness laminates [94]. The model was constructed in ABAQUS and consisted of a 300 mm $\times$ 300 mm plate loaded under compression until failure. The authors found that the tow drop areas, or the gap and overlap locations, were sites for damage initiation. They also stated that the effects of the gaps and overlaps can be incorporated into the design by varying the material allowables. The authors also found that staggering, or dispersing the gaps and overlaps throughout the laminate, improves the overall strength. This model was only theoretical and was not validated with a testing campaign.

Fayazbakhsh et al. built on the model by Blom et al. by establishing a defect layer method to capture the effects of gaps and overlaps [164]. A ply with a gap was modified via its material properties based on a gap area percentage between 0 % (defect-free layer) and 100 % (layer completely made of resin). The variation of the material properties was found via a FE simulation with the results shown in Figure 2.7. The overlap-modified plies maintained the same material properties, but the thickness was increased. Using this method, two designs of variable stiffness cylinders were analyzed. Both analyses showed a reduction in stiffness (-14 % and -15.1 %) and strength (-12.2 % and -12.4 %) when only gaps were addressed. The opposite was true when only overlaps were included, with stiffness increasing by 11 % and 9.6 % respectively for both designs, and strength increasing by 30.5 % and 29.9 %.

Ghayour used a multiscale induced defect layer method, building on the defect layer method from Fayazbakhsh et al., to characterize tow-to-tow gaps in AFP manufactured

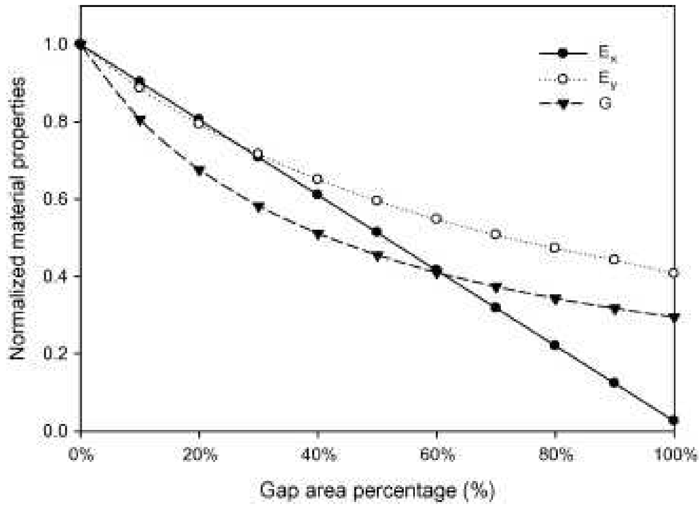

Figure 2.7: Normalized elastic properties with respect to gap area percentage of a given ply [164].

laminates [165]. Gaps within a laminate were mapped back to the mesh in the FE model to establish the gap area within each cell. The results from this analysis were compared with experimental results from Nguyen et al. [160] and showed good agreement with reality.

2.3.5.4 ANGLE DEVIATION AND STEERING DEFECT MODELING

Angle deviation defects are typically not modeled extensively but are extracted from tow and course geometry and relayed back to the design to assess their effects. The deviation between the actual fiber angle ($\mathbf{f}_i$) and the desired angle ($\mathbf{e}_i$) is caused by a combination of tool surface geometry and process planning selections. The value of angle deviation can be found via the vector dot product:

$$\theta_i = \arccos\left(\frac{\mathbf{e}_i \cdot \mathbf{f}_i}{|\mathbf{e}_i||\mathbf{f}_i|}\right) \tag{2.3}$$

Out-of-plane defects, mainly wrinkles, have been modeled and investigated experimentally. Such defects are typically caused by steering in which an in-plane curvature is induced along a course. If a critical limit is reached, out-of-plane defects can occur due to buckling of the tow. The steering radius can be calculated from the geodesic curvature defined by:

$$\kappa_g = \frac{\gamma''(t) \cdot (\mathbf{N} \times \gamma'(t))}{|\gamma'(t)|^3} \tag{2.4}$$

Here, $\gamma(t)$ is the parametric definition of the course or tow centerline and $\mathbf{N}$ is the normal vector to the surface at the analysis point. Wehbe et al. developed a geometrical model for predicting wrinkle formation on tow surfaces [157]. The model developed a parametric representation of the tow surfaces with the rotation angle of the wrinkle approximated by a cosine function. Additionally, wrinkle wavelengths were determined through material properties as formulated by Belhaj and Hojjati, though this requires elaborate material characterization [158]. The model was further improved by incorporating experimental data associated with the wavelength of wrinkles. Yi et al. developed simulative capabilities to study wrinkle formation based on the Rayleigh-Ritz method and minimum potential energy principle [166]. Normal and tangential contact models between the roller and substrate were also developed to improve the accuracy of the wrinkling model. Experimental trials were conducted to verify model accuracy and investigate the influence of various process parameters on wrinkling behavior.

2.3.5.5 IMPLICATIONS FOR STRUCTURAL ANALYSIS

All the defects presented above relate back to the strength of the final product manufactured via AFP. Due to the significance of the effects of defects on overall strength, it is beneficial to briefly review composite failure theories as well. Composite materials display a variety of failure mechanisms due to their complex structure and manufacturing processes. The failure of composites is also multiscale and interactive, with scales ranging from the individual fiber level to the entire structure [167]. These mechanisms include fiber failure, matrix cracking, buckling, and delamination [168]. Some of the most common failure criteria include max stress, Tsai-Wu, Tsai-Hill, and Hoffman [169]. The authors in [168] provide extensive failure criteria equations for each failure mechanism of composites. By default, these failure criteria do not consider the manufacturing characteristics (defects) that are induced during AFP manufacturing. The authors in [136, 160, 170, 171] present methods for characterizing failure mechanisms and predicting failure loads with incorporation of AFP defects. These methodologies allow for more accurate predictions of final part strength when produced with AFP.

With an understanding of how defects form and affect structural performance, the next section examines the optimization efforts that have been developed to minimize defect occurrence and improve overall AFP process efficiency.

2.4 OPTIMIZATION EFFORTS IN AFP

With a foundational understanding of AFP process planning fundamentals and the formation of manufacturing defects, this section reviews the optimization techniques developed to improve the process. Although numerous algorithms have been developed, these efforts have historically been compartmentalized, with optimization often applied to a single aspect of the lifecycle without a holistic view. This approach has led to specialized solutions for part design, manufacturing planning, and machine operation that are not interconnected. To build a comprehensive picture of the state of the art, this review categorizes these efforts into their respective domains. The discussion covers optimizations in part design, the generation of toolpaths at both the ply and laminate level, the tuning of machine process parameters, and nascent efforts in DFM that aim to integrate these traditionally separate activities.

2.4.1 DESIGN OPTIMIZATION

Design optimization in composite manufacturing is a computationally intensive phase that fundamentally differs from designing with isotropic materials like metals. While metal design optimization primarily focuses on part geometry, composite design involves tailoring the part's ply bounds and stacking sequence to achieve specific mechanical performance for a given set of load conditions, typically to minimize weight or maximize stiffness. With this being the most studied field of composites manufacturing, there is a large amount of literature on this topic, and thus a complete review of design optimization is not feasible within the scope of this review since design optimization is not the focus of this dissertation. However, an overview of the work in this field is covered here for background. These optimization activities are often performed in isolation, prior to manufacturing considerations, using FEA tools to evaluate design iterations. This encompasses sizing, stacking sequence optimization, and fiber angle optimization for variable stiffness laminates. These efforts can be broadly delineated into two primary approaches: constant stiffness and variable stiffness optimization.

Conventional composite design relies on constant stiffness laminates, which utilize a uniform stacking sequence, such as a quasi-isotropic layup, across the entire part. The associated optimization problem then becomes finding the optimal stacking sequence of discrete fiber angles (e.g., 0°, 90°, ±45°) that best meets performance criteria. As detailed by Ghiasi et al. [13] and summarized in a review by Nikbakt et al. [172], this process involves finding the ideal combination of critical parameters such as fiber angles, ply thickness, and the number of layers. These optimization efforts are typically applied to common structures like beams, plates, and shells, with the objective of satisfying performance criteria, often based on specific failure theories. While this approach simplifies both design and manufacturing, it does not fully leverage the anisotropic nature of composite materials, as the properties remain constant across the structure's domain.

The advent of AFP, with its ability to continuously steer fibers along a path, has enabled the practical application of variable stiffness designs. This approach offers a significantly expanded design space where the fiber orientation is varied spatially within each ply, allowing stiffness and strength to be tailored directly to local load paths [173]. Research has shown that variable stiffness panels can achieve significant improvements in buckling performance and allow for load redistribution, leading to more efficient and lightweight structures [83, 174–176]. This method, however, requires sophisticated path planning algorithms and advanced optimization routines to define the complex curvilinear fiber trajectories.

A significant limitation of these design optimization methods is that they generally operate under the idealized assumption that the designed part will be a pristinely manufactured part, without defects. They focus on optimizing the theoretical performance of the composite structure but do not account for the manufacturability constraints or the defects that arise during the AFP process. However, as discussed previously, the very geometric complexity introduced by optimized toolpaths, especially in variable stiffness designs, is a primary source of manufacturing defects like gaps, overlaps, and steering violations. This disconnect between the as-designed ideal and the as-manufactured reality means that the full benefits of design optimization may not be realized, often forcing designers to rely on conservative knockdown factors to account for manufacturing uncertainties.

2.4.2 PATH PLANNING OPTIMIZATION

Path planning optimization is the geometric component of process planning, focused on defining the specific toolpaths the machine will follow to manufacture the structure. It is one of the most difficult aspects of AFP, involving the selection of numerous parameters to generate paths for a given geometry. The heart of this task is selecting appropriate layup strategies and starting points for each ply in the laminate. Historically a manual, trial-and-error process reliant on expert knowledge, recent research has focused on developing algorithms to automate and optimize these critical decisions. These efforts can be divided into two distinct but related challenges: ply-level and laminate-level optimization.

2.4.2.1 PLY-LEVEL OPTIMIZATION

PLO focuses on optimizing the manufacturing strategy for a single ply to minimize predictable, geometry-based defects such as gaps, overlaps, angle deviation, and steering violations. To automate this task, a manufacturability score, or objective function, is required to numerically quantify the quality of a given ply scenario (a specific combination of a starting point and layup strategy). Harik et al. [177] developed the CAPP software to accomplish this, using defect predictions from VCP to score scenarios. Since evaluating this objective function requires computationally expensive VCP simulations for each potential scenario, surrogate-based methods like Bayesian Optimization have been shown to be effective for efficiently finding optimal inputs [75, 178]. Other research has focused on developing path generation methods to specifically manage the trade-off between angle deviation and wrinkle formation on complex surfaces [179], as well as preceding this with a placement suitability analysis to select appropriate parameters and reduce iterations [180]. Additional work by Debout et al. [181] has focused on optimizing the kinematics of a predefined toolpath, using least-squares algorithms to smooth machine motion and reduce manufacturing time.

2.4.2.2 LAMINATE-LEVEL OPTIMIZATION

While PLO is effective for enhancing individual plies, an optimal set of individual plies does not guarantee an optimal laminate. The more complex challenge of LLO addresses the through-thickness interaction of defects between plies. The primary objective of LLO is to minimize defect stacking and promote the dispersion of defects throughout the laminate, a practice commonly known as staggering. This becomes a difficult combinatorial optimization problem, as a laminate of n plies with m possible scenarios for each ply results in m^n total combinations, making a brute-force search of the design space computationally infeasible.

To solve this, researchers have investigated various combinatorial optimization algorithms. As explored by Swingle [76] and Patel [77], a simple Greedy Search (GS) algorithm can be used to provide a rapid solution by making a locally optimal choice at each ply. However, this method rarely finds the globally optimal laminate. To better explore the vast design space, metaheuristic algorithms that incorporate a stochastic search process have been implemented. These population-based methods, including Genetic Algorithms (GA), Differential Evolution (DE), and Particle Swarm Optimization (PSO), have been shown to be more effective at identifying globally near-optimal laminate configurations that minimize defect stacking [182, 183]. These advanced algorithms provide a robust method for achieving intelligent staggering, not just between plies of the same angle, but across the entire laminate structure.

2.4.3 PROCESS PARAMETER OPTIMIZATION

Beyond the geometric generation of toolpaths and structural optimizations, the quality of an AFP-manufactured part is profoundly influenced by the physical machine settings used during layup. Process parameter optimization focuses on identifying the ideal combination of these settings, such as temperature, compaction pressure, layup speed, and tow tension, to achieve desired outcomes. The primary objective is often to balance the competing demands of maximizing production throughput while ensuring adequate material consolidation, interlaminar bond quality, and the prevention of defects. Research in this area has utilized a range of methods, from physics-based models and statistical analysis to more recent data-driven machine learning techniques.

Initial efforts in this domain often focused on the intricate interplay between multiple parameters. Aized and Shirinzadeh [184] utilized a response surface method to investigate the influence of gas torch temperature, head speed, and compaction force, revealing an optimal temperature threshold to avoid thermal degradation. Their work also confirmed the inverse relationship between layup speed and quality, which they attributed to reduced dwell time for intimate heating and contact. Similarly, Yu et al. [185] highlighted parameter sensitivity in the related tape winding process, defining optimal ranges for temperature, tension, pressure, and velocity that yielded superior interlaminar shear strength.

More advanced studies have employed multi-scale modeling to capture complex physical interactions with greater fidelity. For instance, Sun et al. [186, 187] demonstrated a collaborative optimization approach using a combination of FE and Molecular Dynamics (MD) methods to improve comprehensive multiscale mechanical performance. Taking a different approach, Han et al. [188] proposed a low-entropy design method that uses enthalpy and entropy calculations as an indirect but effective measure of quality. By establishing a link between the absorption of energy within the AFP system and a reduction in void content, their work augmented the mechanical properties of the final laminates. In the domain of thermoplastic composites, Sonmez and Akbulut [189] developed a process optimization scheme focused on heating parameters and roller velocity to minimize peak tensile residual stress, concluding that a compromise on layup speed was necessary for achieving optimal quality.

In addition to multi-parameter studies, other research provides a more nuanced understanding by focusing on the optimization of a single, critical parameter. Several investigations have centered on compaction pressure, developing FE models to analyze the adaptability of rollers to molds and optimize their design for uniform pressure distribution, which is critical for preventing defects and improving bond strength [190, 191]. For laser-heated systems, Maurer and Mitschang [192] identified critical process temperatures and demonstrated that diode lasers offer remarkable efficiency for thermoplastic tape placement. Tow tension has also been isolated, with dedicated studies developing and validating equations to capture tension variations, leading to the identification of optimal tension levels and improved designs for tow feeding systems [191, 193].

Complementing these physics-based and statistical techniques, recent advancements have introduced ML as a powerful tool for process parameter optimization. These data-driven methods are particularly effective where process dynamics are highly non-linear. In a study by Bruning et al. [57], a machine learning approach was used to plan, optimize, and inspect the AFP process by aggregating vast amounts of data from planning, CNC operations, and real-time monitoring. This work also introduced a closed-loop optimization concept, where data captured during the physical layup process is used to dynamically inform subsequent parameter adjustments. Similarly, Heider et al. [194] presented an on-line optimization algorithm using Artificial Neural Networks (ANNs) for an automated thermoplastic placement system. The ANNs-based models were designed to predict material quality from various process set points, enabling the system to maximize throughput while maintaining a specified quality level. These studies demonstrate the potential for ML to bridge the gap between real-time process data and actionable optimization decisions, harnessing numerical optimization to balance quality specifications with throughput demands.

2.4.4 HARDWARE AND KINEMATIC OPTIMIZATION

AFP has witnessed continuous advancements in its hardware components to enhance the reliability, efficiency, and quality of fiber placement. Addressing concerns ranging from tension control to toolpath design, various innovative solutions have been introduced. Several notable hardware optimization efforts have been reported in the literature. One of the revolutionary developments in AFP hardware is the modular-servo-creel head [58]. Designed to address the reliability issues observed during the production of the 777x spar, this hardware component offers solutions to loss of tension that typically occurs during substantial speed changes, especially during zero-degree plies. The introduction of an

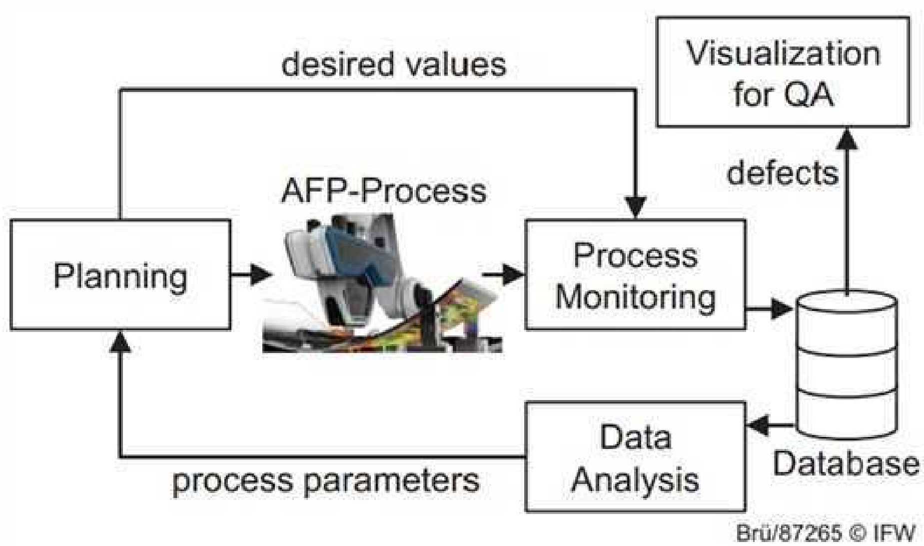

Figure 2.8: Approach of data aggregation-based process optimization [57].

improved control architecture in this head significantly mitigates tow placement errors arising from acceleration and deceleration during feed and cut cycles.

The servo-servo creel system presents a significant shift from traditional tension control mechanisms [135]. Replacing the conventional brake disc, the servo-servo creel incorporates a servo/controller and gearbox combination. This shift to motor controller-based tension control instead of PLC-based methods enhances sensitivity and reaction time. The system is designed to support high-speed layups and deliver finer tension control essential for dry-fiber placement, where material lacks resin to hold the fibers together. With its improved sensitivity, the servo-servo creel can react promptly to provide reduced operating tension at higher speeds, crucial for preventing shearing and undesired movement of fibers in the substrate. Moreover, the integration of internal sensors under the RIPITx banner enables the servo-servo creel to measure issues like slipped tows, add/cut error detection, and tow-end placement.

Debout et al. [181] put forward a methodology to elevate the efficiency and quality of AFP machines by focusing on toolpath smoothing. The method, aimed at resolving machine redundancy and refining the toolpath, minimizes variations in rotational axis control values. This optimization enhances the kinematic behavior, thus reducing manufacturing time. Validated on a model of an airplane nose, the results shown in Figure 2.9 indicate that this method, based on the tool axis adherence to the digital admissible orientation, upholds the desired quality standards.

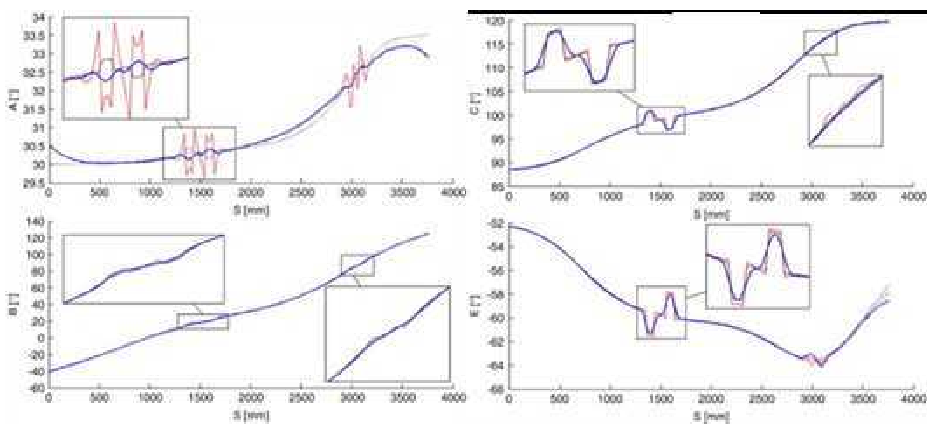

Figure 2.9: A, B, C, E axis position orders for original (dotted line), smoothed optimized (blue line) and non-smoothed optimized toolpaths (red line) [181].

Zhang et al. [134] introduced a methodological approach to AFP mechanism layout design, which involves breaking down the mechanism into four subsystems. By applying coordinate transformation, the orientation and position of each subsystem are articulated as design variables. The optimization strategy, illustrated in Figure 2.10, endeavors to minimize the space occupied, reduce fiber defects, and mitigate bending during the transmission path. The constraints set in place ensure zero interference between adjacent objects and the fiber path. The verification results, conducted on 4-strip and 6-strip AFP mechanisms, confirm the effectiveness of this approach, holding promise in minimizing the physical damage to prepreg tow and internal defects in preforms.

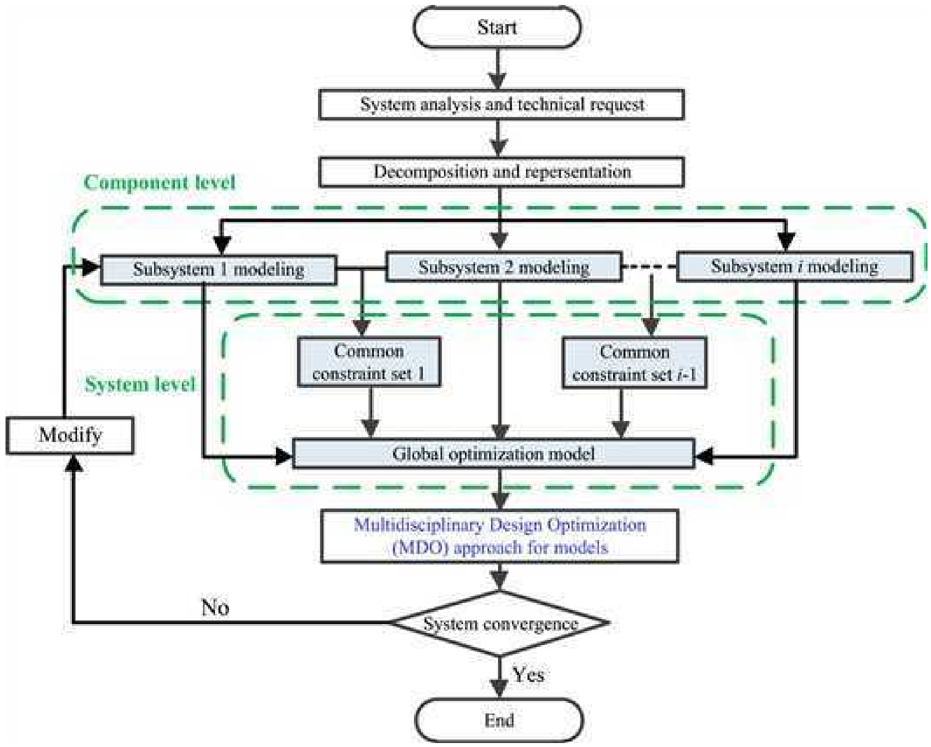

Figure 2.10: Layout optimization process for AFP mechanisms [134].

2.4.5 DESIGN FOR MANUFACTURING

While the previously discussed optimization efforts often operate within isolated domains, a growing body of research has focused on DFM principles that seek to integrate the traditionally separate stages of the AFP lifecycle. Even though translating a design into a manufacturing plan can have large effects on the resulting structure, limited research has been published on developing standardized methods for this purpose [7, 53]. Such an effort requires close cooperation between designers, process planners, and manufacturing personnel to ensure decisions made during the design phase account for manufacturability and that manufacturing-induced defects are fed back to inform the final performance of the part. A core enabler of this integration is the development of robust data mapping techniques. As detailed by Brasington et al., this work establishes a uniform mapping framework that ensures the interoperability of multimodal data, spanning from numerical analysis to in-situ inspection, by relating these various conditions to a specific position on the tool surface [195].

A primary focus of DFM in AFP has been the development of workflows that, while powerful, often rely on manual data exchange to map predicted manufacturing characteristics back to structural analysis models. Initial work in this area by Noevere et al. [196–199] established a communication channel between design optimization software, such as Collier Aerospace's HyperX, and composite programming software, like VCP. An example of this mapping methodology on a "saddle" shape is shown in Figure 2.11. This methodology allows manufacturing-induced defects predicted by VCP, such as fiber angle deviation, gaps, and overlaps, to be mapped back onto the FEM used for design analysis. With this data integrated, designers can re-run the analysis to evaluate safety margins and update ply boundaries or stacking sequences based on a more realistic, as-manufactured state rather than an idealized one.

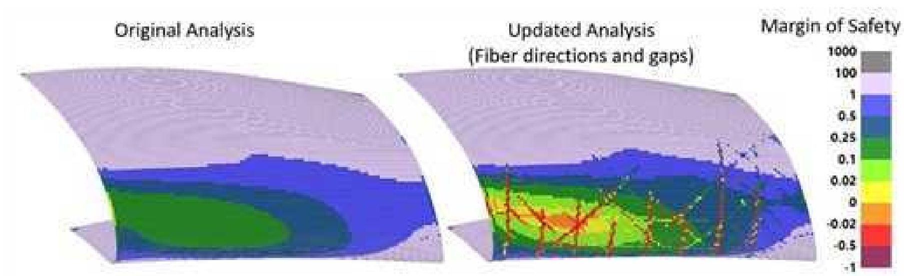

Figure 2.11: Example defect mapping for angle deviation and gaps [196].

Building on this mapping capability, a crucial advancement was the development of a methodology to directly incorporate design data into process planning metrics in an automated fashion. As detailed by Brasington et al. [200], this is achieved by importing safety margins from HyperX into the CAPP software. Within CAPP, the predicted AFP defects are mapped to these margin values, and the resulting, updated margins are then incorporated into the manufacturability scoring algorithms. This creates a design-informed process planning analysis where the optimization objective itself is influenced by the structural importance of different regions of the part.

These DFM concepts were enabled by the development of dedicated software tools. Halbritter developed an initial CAPP software prototype to semi-automate the optimization of essential process planning inputs, namely layup strategy and starting point selection [177]. In a parallel effort to close the design-analysis-manufacturing loop, personnel from Boeing and Vericut developed a "laminate manager" within VCP [201]. This tool provides both designers and process planners with a laminate-based focus, allowing them to visualize and analyze each ply in the context of the full laminate. By enabling both teams to work from a shared understanding of defect behavior, these tools aim to reduce development timelines and streamline the certification process.

The optimization efforts reviewed in this section, while individually effective, remain largely compartmentalized within their respective domains. The following section synthesizes the literature to identify the critical gaps that persist and motivate the integrated approach presented in this dissertation.

2.5 SYNTHESIS AND GAPS

The preceding sections provide a comprehensive view of the AFP lifecycle, detailing the fundamental principles of design, the intricacies of path planning, the mechanisms of defect formation, and the various efforts to optimize the process. The research demonstrates a mature understanding of individual domains: FEA is well-established for design optimization, numerous algorithms exist for optimizing toolpaths, and a variety of models have been developed to predict the influence of process parameters. Recent work in DFM further illustrates a clear trend toward integrating these disparate stages, primarily through the mapping of predicted manufacturing data back to design models. However, despite these advancements, the prevailing approach remains fragmented. The optimization methods are rarely interconnected, and a truly holistic, closed-loop system that spans the entire lifecycle has not been realized. Based on the literature reviewed, several critical gaps can be identified:

1. **Open-Loop and Siloed Pillars:** The most significant gap is the persistent open-loop nature of both modeling and optimization in AFP. Design, process planning, manufacturing, and inspection are still treated as largely separate activities with limited, often manual, feedback between them. There is a distinct absence of a holistic, interconnected system that optimizes the process as a whole and allows data and insights from one stage to dynamically influence decisions in another.

2. **Decoupled Structural and Manufacturing Optimization Metrics:** While methods exist to map defect predictions back to structural models, the reverse is less developed. Process planning optimization is often performed without consideration of the structural criticality of different regions of a part. Defect scoring and optimization algorithms typically treat all defects of a certain type and size equally, without considering whether a defect is located in a high-margin area or a region critical to the component's safety and performance.

3. **Inefficient and Non-Integrated Optimization Workflows:** Current process planning relies heavily on the experience of engineers performing manual, iterative tasks. While initial research prototypes have begun to explore the semi-automation of ply-level decisions, these efforts remain limited in scope and are not integrated with laminate-level strategies, structural feedback, or inspection data. There is therefore a need for a comprehensive, purpose-built optimization framework that addresses these limitations systematically. The reliance on external, computationally expensive software for every iteration of a toolpath evaluation creates a bottleneck, slowing down the exploration of the design space and hindering rapid optimization.

4. **Absence of As-Manufactured Data in the Feedback Loop:** Existing DFM workflows

primarily focus on a one-way transfer of information from process planning (predicted defects) back to design. The loop is rarely closed with data from the final, as-manufactured part. Without integrating post-production inspection data, the predictive models for defect formation cannot be validated or refined, and the system cannot learn from real-world production outcomes to improve future manufacturing cycles.

Collectively, these gaps highlight the need for a smart, data-driven, and lifecycle-integrated approach to AFP process planning. The research presented in this dissertation directly addresses these shortcomings through the development of the iSPP framework. To address the first gap, the iSPP establishes a closed-loop architecture that connects design, process planning, manufacturing, and inspection through bidirectional data exchange. The second gap is addressed through the DIDS, which incorporates structural margin-of-safety data from HyperX into the defect scoring methodology, enabling process planning decisions to be informed by the structural criticality of different regions. The third gap is resolved through the CAPP software and its PLO and LLO algorithms, which provide systematic, semi-automated optimization of layup strategies and starting points at both the individual ply and full laminate levels. Finally, the fourth gap is closed by integrating as-manufactured inspection data from RIPITx into the CAPP environment, enabling comparison of predicted and observed defects to support continuous process refinement. The following chapter presents the methodology of the iSPP framework in detail.

CHAPTER 3

THE INTEGRATED SMART PROCESS PLANNING (iSPP) FRAMEWORK

The literature review presented in Chapter 2 identified several gaps in current AFP process planning practice: the absence of unified metrics that quantify manufacturability across multiple defect types, the reliance on manual and inefficient optimization workflows for process planning decisions, limited integration between structural design and manufacturing planning, and the lack of closed-loop feedback connecting inspection results to process improvement. This chapter presents the iSPP framework, a methodology that addresses these gaps by establishing CAPP as a central integration platform connecting design intent, defect prediction, process optimization, and part inspection.

The iSPP framework operates through bidirectional data exchange with established industrial tools. Structural analysis and laminate optimization are performed in HyperX, which provides element-level MoS values that quantify local structural criticality. Path generation, defect prediction, and Numerically Controlled (NC) code output are performed in VCP, which serves as the simulation kernel for evaluating candidate process plans. Inspection data from RIPITx, an automated in-process inspection system, provides as-manufactured defect measurements that enable comparison with as-designed predictions. CAPP orchestrates these tools while providing scoring algorithms that translate defect predictions into quantitative metrics and optimization routines that search the process planning design space.

This chapter presents each component of the framework in sequence. Section 3.1 describes the subprocess integration with VCP, defining the data exchange protocol and defect classification scheme that underlies subsequent analysis. Section 3.2 introduces the MIDS, a scoring formulation that quantifies manufacturing quality based on defect frequency and severity. Section 3.3 extends this formulation to incorporate structural margins through the DIDS, and introduces the IDS, which combines both perspectives through user-defined weighting to enable trade-off analysis between manufacturing efficiency and structural performance. Section 3.4 presents the PLO algorithm for optimizing individual ply configurations, while Section 3.5 extends optimization to the laminate level through the LLO algorithm. Section 3.6 describes the incorporation of inspection data for as-manufactured comparison. Finally, Section 3.7 synthesizes these components into a unified digital thread connecting design through inspection. Verification and demonstration of each component through systematic application to the Quartic benchmark surface is presented in Chapter 4.

3.1 VERICUT COMPOSITES PROGRAMMING (VCP) INTEGRATION

VCP, developed by Vericut, is an industry-standard software platform for programming and simulating automated composite manufacturing processes such as AFP and ATL [202]. The software provides comprehensive functionality for ply creation, path generation, defect prediction, and NC code output, enabling the translation of manufacturing plans into machine-executable programs. VCP supports a wide range of manufacturing systems and maintains compatibility with major CAD and composite design environments, including CATIA (used in this work). Through its companion module, Vericut Composites Simulation (VCS), users can validate NC programs for potential collisions or path errors, providing a comprehensive virtual representation of composite fabrication.

Within the iSPP framework, VCP functions as the *analysis kernel*, providing validated algorithms for ply generation, course creation, and defect prediction that are accessed programmatically through the CAPP environment. Rather than replicating these mature, industry-accepted functions, CAPP leverages VCP's capabilities to ensure that its results remain consistent with established industrial practice. This design philosophy positions CAPP as an optimization and decision-support layer built upon VCP's simulation foundation. The integration is achieved through VCP's Python Application Programming Interface (API), which allows CAPP to execute VCP operations in a headless manner—without requiring manual interaction with the VCP graphical interface.

This section defines the data exchange protocol that governs communication between CAPP and VCP. The protocol comprises three components: (1) the finite element discretization that provides a common geometric reference, (2) the input parameters that CAPP supplies to VCP for ply and laminate scenarios, and (3) the output data structures through which VCP returns course geometry and defect predictions. These definitions establish the foundation for the defect scoring methodology presented in Section 3.2.

3.1.1 FINITE ELEMENT DISCRETIZATION

The geometric foundation for the CAPP–VCP integration is a finite element mesh that discretizes the tool surface into a set of quadrilateral elements. This mesh serves as the common spatial reference shared among CAPP, VCP, and HyperX, enabling consistent mapping of defect predictions, structural margins, and optimization variables to the same geometric domain.

Let $\mathcal{E} = \{e_1, e_2, \ldots, e_{N_e}\}$ denote the set of finite elements, where N_e is the total number of elements within the ply boundary. Each element e_n is defined by its nodal connectivity (four corner vertices), centroid coordinates $\mathbf{c}_n = (x_n, y_n, z_n)$, and projected surface area A_n. The total ply area is computed as the sum of individual element areas:

$$A_{\text{ply}} = \sum_{n=1}^{N_e} A_n \tag{3.1}$$

The mesh is defined by the user during project initialization and exported to VCP as part of the input data package, as illustrated in Figure 3.1. VCP uses the element centroids as evaluation points for point-based defect metrics, ensuring that defect values are spatially registered to the same elements used for structural analysis in HyperX. This shared discretization is essential for the margin-weighted scoring methodology introduced in Section 3.3, where defect penalties must be correlated with element-level structural margins.

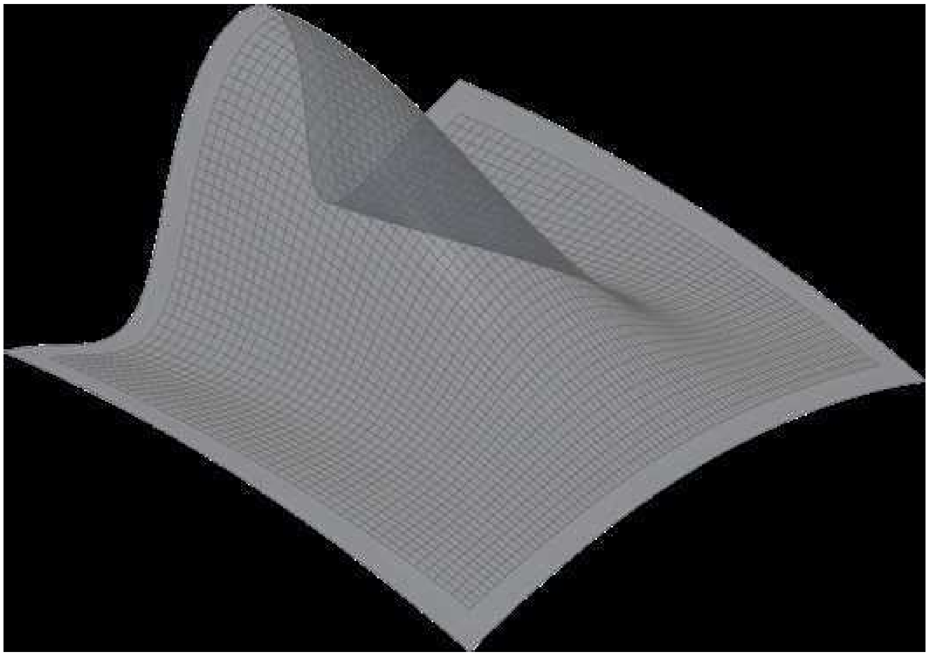

Figure 3.1: Finite element mesh visualization within CAPP, showing the quadrilateral element discretization of the tool surface.

3.1.2 INPUT PARAMETER SPACE

For each ply scenario, CAPP assembles an input parameter set and transmits it to VCP through the subprocess interface. These parameters are organized into three categories: geometric definitions, process planning variables, and coordinate system specifications.

The geometric definitions include the tool surface geometry, laminate definition file, and the finite element mesh described in Section 3.1.1. These files establish the physical domain of the tool surface over which VCP generates courses and evaluates defects. The tool surface is typically imported from a CAD environment as a STEP file, while the laminate definition specifies ply boundaries, stacking sequence, and fiber orientations.

The process planning variables constitute the controllable parameters that CAPP manipulates during optimization. For each ply, these include the starting point coordinates $\mathbf{p}_0 = (x_0, y_0, z_0)$, the number of tows per course, and the layup strategy (e.g., parallel, rosette, or natural path generation). The starting point defines the location from which course propagation initiates, while the layup strategy governs how subsequent courses are generated relative to the reference curve. These variables form the decision space for the PLO and LLO algorithms described in Section 3.4 and Section 3.5, respectively.

The coordinate system specifications define the angular reference axis, axis system orientation, wrap type, and unit system. These settings ensure geometric consistency between the CAD model, VCP's internal computations, and CAPP's visualization and analysis routines. Unit alignment is verified during import to prevent scaling errors that would propagate through the defect scoring calculations.

3.1.3 OUTPUT DATA STRUCTURE

Upon completing ply generation and defect analysis, VCP produces a set of output files that encode course geometry, defect locations, and defect attributes. For each ply, four primary files are generated: a centroid-based defect file, a laser-format geometry file, and two defect attribute files for gaps and overlaps respectively. Table 3.1 summarizes the contents of each file.

Table 3.1: Summary of VCP output files generated for each ply scenario.

File	**Format**	**Contents**
`Ply_N_Centroids.FEA.xml`	XML	Point defect values at element centroids
`Ply_N_LaserEI.laserdata.xml`	XML	Polygon geometries for area defects
`Ply_N-Gaps.csv`	CSV	Gap attributes (length, width, area)
`Ply_N-Overlaps.csv`	CSV	Overlap attributes (length, width, area)

The centroid file contains one record for each element in the finite element mesh. Each record includes the element identifier, centroid coordinates, and, where applicable, defect values for angle deviation, steering radius, and roller compression. Summary statistics for gaps and overlaps, including total area, minimum and maximum values, and stackup counts, are provided in the file header, but the detailed polygon geometries required for element intersection calculations are contained in the laser-format geometry file. The element identifiers in this file correspond directly to those in the input mesh, enabling immediate spatial registration without coordinate transformation.

The laser-format geometry file provides detailed polyline representations of ply boundaries, course centerlines (headpaths), tow edges, and defect polygon outlines. Gap and overlap defects are encoded as closed polylines, with each vertex defined by its spatial coordinates and surface normal vector. This geometric representation enables visualization within CAPP and provides the polygon boundaries required for element intersection calculations.

The gap and overlap attribute files contain summary statistics for each discrete defect instance. Each record includes a defect identifier, total length, maximum width, and enclosed area. The maximum width attribute serves as the basis for threshold comparisons in the scoring methodology, while the polygon geometry from the laser file enables spatial mapping to individual elements.

3.1.4 DEFECT CLASSIFICATION AND ATTRIBUTES

The defects predicted by VCP are classified into two categories based on their geometric representation: point defects and area defects. Let $\mathcal{M}$ denote the set of all defect types considered in the scoring framework:

$$\begin{aligned} \mathcal{M} &= \mathcal{M}_{\text{point}} \cup \mathcal{M}_{\text{area}} \\ \mathcal{M}_{\text{point}} &= \{\text{angle deviation}, \text{steering}\} \\ \mathcal{M}_{\text{area}} &= \{\text{gap}, \text{overlap}\} \end{aligned} \tag{3.2}$$

3.1.4.1 POINT DEFECTS

Point defects are evaluated at element centroids and characterized by a scalar magnitude. *Angle deviation* quantifies the angular difference, in degrees, between the as-planned fiber orientation and the as-manufactured fiber direction resulting from course geometry and surface curvature. *Steering* quantifies the in-plane curvature of the tow path, reported by VCP as a steering radius in length units. Because tighter radii correspond to more severe manufacturing risk, the steering value is internally converted to curvature (the reciprocal of radius) for threshold comparison, such that larger values indicate greater severity. The curvature κ_n at element n is computed from the VCP-reported steering radius r_n as:

$$\kappa_n = \frac{1}{r_n} \tag{3.3}$$

with $\kappa_n = 0$ assigned to elements where the path is straight ($r_n \to \infty$).

For each point defect type $m \in \mathcal{M}_{\text{point}}$, the defect value at element n is denoted $d_{m,n}$. Elements where VCP reports no defect, or where the tow path is straight (infinite steering radius), are assigned a defect value of zero.

3.1.4.2 AREA DEFECTS

Area defects arise from geometric interactions between adjacent courses and are represented as closed polygons. *Gaps* occur where insufficient material coverage leaves voids between neighboring tow edges, while *overlaps* occur where excessive material coverage causes adjacent tows to superimpose. Each area defect instance j is characterized by three attributes: total length ℓ_j, maximum width w_j, and enclosed area A_j^{defect}.

The maximum width w_j serves as the primary attribute for threshold comparison, consistent with industry practice where gap and overlap tolerances are typically specified as width limits. The polygon geometry defines the spatial extent of the defect and is used to compute the intersection with individual finite elements, as described in Section 3.1.5.

3.1.5 SPATIAL REGISTRATION

To enable element-level scoring, defect data from VCP must be mapped onto the finite element mesh. The mapping procedure differs for point and area defects, reflecting their distinct geometric representations.

For point defects, the mapping is direct. The centroid file produced by VCP uses element identifiers that correspond exactly to those in the input mesh. When CAPP imports the centroid file, each defect value $d_{m,n}$ is assigned to element e_n by identifier lookup, with no geometric computation required. This correspondence is guaranteed by VCP's use of the user-supplied centroid locations as evaluation points.

For area defects, the mapping requires computation of polygon-element intersections. Let D_j denote the polygon boundary of area defect j, and let e_n denote element n with area A_n. The coverage fraction $\alpha_{j,n}$ quantifies the proportion of element n that is covered by defect j:

$$\alpha_{j,n} = \frac{\mathrm{Area}(D_j \cap e_n)}{A_n} \tag{3.4}$$

where $\mathrm{Area}(D_j \cap e_n)$ is the area of intersection between the defect polygon and the element. The coverage fraction satisfies $0 \leq \alpha_{j,n} \leq 1$, with $\alpha_{j,n} = 0$ indicating no intersection and $\alpha_{j,n} = 1$ indicating complete coverage.

The intersection calculation is implemented using a two-stage algorithm. First, a k-d tree is constructed from the element centroid coordinates, and nearest-neighbor queries identify candidate elements in the vicinity of each defect polygon vertex. This spatial indexing reduces the number of intersection tests from $O(N_e)$ to approximately $O(k)$ per defect, where $k \ll N_e$ is the number of neighboring elements. Second, for each candidate element, the polygon intersection area is computed using computational geometry routines operating on the two-dimensional parametric domain of the tool surface. The resulting coverage fractions are stored in an element-indexed data structure for subsequent use in the penalty calculations.

Once VCP completes ply generation and CAPP imports the resulting defect data, the mapped defect values are visualized directly on the tool surface to support rapid assessment (Figure 3.2). Area defects are rendered as closed polygon features within the ply boundary, with gaps shown in green and overlaps shown in red. Course centerlines are shown as dashed lines to enable spatial correlation between toolpath geometry and predicted defect locations. Point defects—angle deviation and steering—are displayed as a color-mapped point cloud overlaid on the element mesh, where color intensity corresponds to defect magnitude at each element centroid.

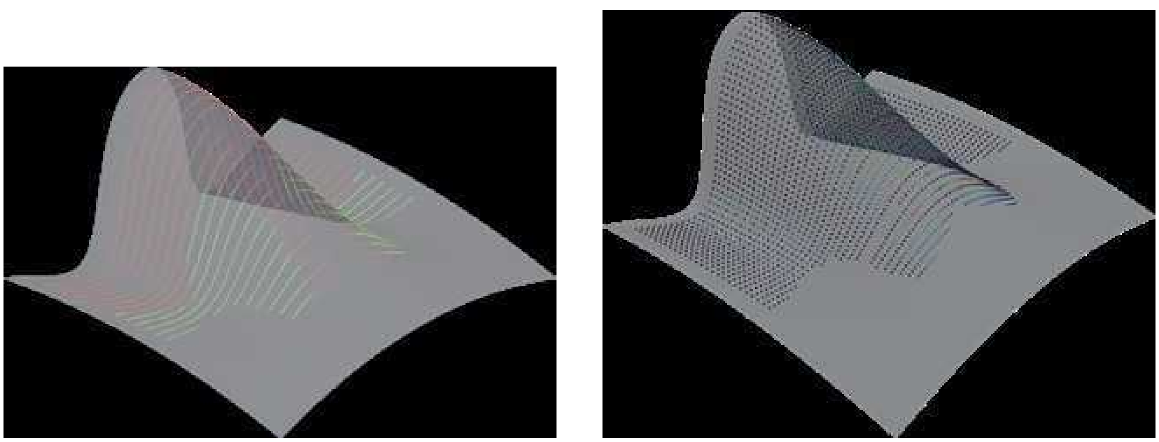

(a) Area defects: course centerlines (dashed light blue), gaps (green), and overlaps (red).

(b) Point defects: severity displayed as a color-mapped point cloud on the element mesh.

Figure 3.2: Visualization of VCP output within CAPP. Area and point defect layers are displayed independently to support targeted assessment of each defect category.

3.1.6 SUBPROCESS ARCHITECTURE

Early versions of the CAPP–VCP integration employed a socket-based network protocol to maintain communication between the two applications [203]. Although functional, this method suffered from limited portability, fragile network dependencies, and minimal user feedback, frequently leading to uncertain execution states outside controlled environments. These limitations motivated a ground-up redesign of the integration layer into an enhanced subprocess-based architecture designed to provide stable, high-fidelity communication.

In the current implementation, CAPP operates as the controlling process that programmatically spawns an independent worker subprocess executing VCP's Python API. The controlling process manages the full subprocess lifecycle, including initialization, data transfer, execution monitoring, and controlled termination. Communication is implemented through Inter-Process Communication (IPC) channels that transmit commands, status messages, and error codes in a non-blocking manner. This architecture isolates VCP execution from the primary CAPP application thread, preventing interface freezes during intensive operations and ensuring that unhandled exceptions within VCP do not crash the overall application.

From the user's perspective, the integration functions as a single application. During execution, VCP console outputs are streamed directly into CAPP's log, and a dedicated progress indicator displays the current ply index and total count (Figure 3.3a). Error-handling routines trap and report exceptions with diagnostic information and suggested corrective actions (Figure 3.3b).

A key advancement of this architecture is its support for ply-level reprocessing. The socket-based method required complete laminate regeneration for every parameter change: a new laminate file was created, all process parameters were redefined, and every ply was regenerated serially. In contrast, the subprocess integration allows CAPP to load a

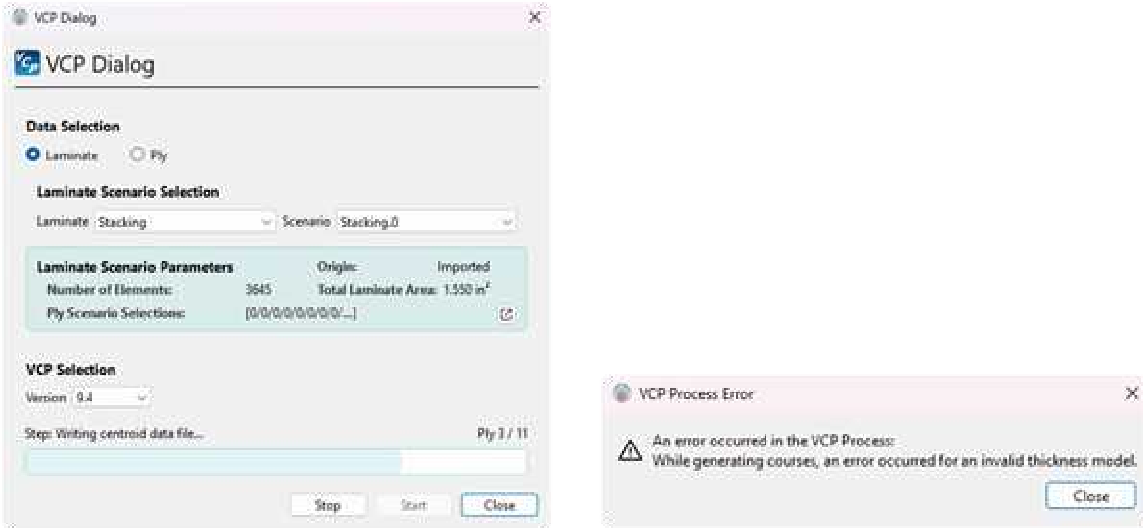

(a) (b)

Figure 3.3: (a) VCP subprocess dialog displayed within CAPP and (b) an example error message generated during the VCP process.

pre-existing `.VCLaminate` file and modify only the specific ply under consideration. When an optimization routine adjusts a starting point or layup strategy, the subprocess instructs VCP to regenerate only that ply while leaving others untouched. This granular control is enabled by the persistent VCP process, which maintains the laminate data structure in memory between calls. The functional correctness and efficiency characteristics of this architecture are demonstrated through quantitative verification in Chapter 4.

3.2 MANUFACTURING-INFORMED DEFECT SCORE (MIDS)

The defect predictions produced by VCP provide detailed spatial information about gaps, overlaps, angle deviations, and steering violations across the ply domain. While this information is valuable for identifying problem areas, it does not directly answer the process planner's central question: *how manufacturable is this ply scenario?* Comparing two candidate scenarios requires synthesizing multiple defect types, each with different units and physical interpretations, into a coherent assessment.

The MIDS addresses this need by transforming the element-level defect data into a single scalar metric bounded between 0 and 100. This formulation provides an intuitive interpretation analogous to a percentage score: a MIDS of 100 indicates that no defects exceed their respective thresholds (ideal manufacturability), while lower scores reflect increasing penalty contributions from threshold-exceeding defects. The score's bounded range enables direct comparison across ply scenarios, layup strategies, and starting point configurations without requiring the user to interpret raw defect counts or areas.

This section presents the mathematical formulation of MIDS, beginning with the defect criteria and weight assignment methodology, proceeding through the threshold, severity, and penalty functions, and concluding with the score aggregation that produces the final metric. The formulation operates at the ply level, evaluating manufacturability for a single ply scenario; extension to margin-weighted scoring is presented in Section 3.3. Demonstration of the MIDS discriminative capability on the Quartic benchmark surface is presented in Section 4.3.

3.2.1 DEFECT CRITERIA AND WEIGHT ASSIGNMENT

The MIDS formulation requires two inputs for each defect type $m \in \mathcal{M}$: a threshold τ_m that defines the boundary between acceptable and penalized defects, and a pair of weights ω_m^{freq} and ω_m^{sev} that control the relative importance of that defect type in the overall score.

3.2.1.1 THRESHOLD DEFINITION

The threshold τ_m establishes the minimum defect magnitude that contributes to the penalty. Defects below the threshold are considered acceptable and do not affect the score. Table 3.2 summarizes the threshold interpretation for each defect type.

For gap and overlap defects, the threshold is compared against the maximum width of the defect polygon, consistent with industry practice where coverage tolerances are

Table 3.2: Defect threshold definitions and units.

Defect Type	**Threshold Attribute**	**Units**	**Interpretation**
Gap	Maximum width	Length	Widest point of gap polygon
Overlap	Maximum width	Length	Widest point of overlap polygon
Angle deviation	Fiber misalignment	Degrees	Deviation from nominal orientation
Steering	Path curvature[a]	1/Length	In-plane curvature of tow path

[a] User specifies minimum steering radius; internally converted to maximum curvature.

specified as width limits. For angle deviation, the threshold is compared directly against the angular difference in degrees. For steering, the user specifies a minimum allowable radius (a more intuitive quantity), which is internally converted to a maximum allowable curvature $\tau_{\text{steering}} = 1/r_{\text{min}}$ for threshold comparison.

3.2.1.2 WEIGHT ASSIGNMENT VIA ANALYTIC HIERARCHY PROCESS

The weights ω_m^{freq} and ω_m^{sev} determine how each defect type contributes to the total penalty. These weights are assigned using the Analytic Hierarchy Process (AHP), a structured decision-making methodology that converts pairwise comparisons of criteria importance into a consistent set of numerical weights [204].

In the AHP framework, the user constructs a pairwise comparison matrix where each entry reflects the relative importance of one defect type versus another. For example, if gap defects are judged to be moderately more important than steering defects for a particular application, this judgment is encoded as a numerical ratio. The principal eigenvector of the comparison matrix yields normalized weights that satisfy the consistency requirements of the method.

For each defect type, two weights are assigned: ω_m^{freq} governs the penalty contribution from defect frequency (the fraction of ply area affected), while ω_m^{sev} governs the contribution from defect severity (the magnitude of threshold exceedance). The combined weight for defect type m is:

$$\omega_m^{\text{combined}} = \omega_m^{\text{freq}} + \omega_m^{\text{sev}} \tag{3.5}$$

The weights are subject to a normalization constraint ensuring the total penalty is bounded:

$$\sum_{m \in \mathcal{M}} \omega_m^{\text{combined}} = \sum_{m \in \mathcal{M}} \left(\omega_m^{\text{freq}} + \omega_m^{\text{sev}} \right) = 100 \tag{3.6}$$

This constraint guarantees that the MIDS remains in the interval $[0, 100]$ regardless of the specific weight distribution. A representative weight configuration is shown in Table 3.3, where equal importance is assigned to all defect types and equal emphasis is placed on frequency and severity within each type.

Table 3.3: Example defect criteria configuration with equal weighting.

Defect Type	τ_m	**Units**	ω_m^{freq}	ω_m^{sev}	$\omega_m^{\text{combined}}$
Gap	0.125	in.	12.5	12.5	25
Overlap	0.125	in.	12.5	12.5	25
Angle deviation	5.0	deg	12.5	12.5	25
Steering	15.0	in.[a]	12.5	12.5	25
Total			50	50	100

[a]Specified as minimum radius; converted to curvature threshold $\tau = 1/15.0 = 0.0667$ in.$^{-1}$.

3.2.2 THRESHOLD INDICATOR FUNCTION

The threshold indicator function determines whether a defect contributes to the penalty calculation. Defects that do not exceed their respective thresholds are excluded from the score, reflecting the practical reality that small deviations within tolerance do not constitute manufacturing defects.

For point defects (angle deviation and steering), the indicator function evaluates the defect value $d_{m,n}$ at element n against the threshold τ_m:

$$\mathbf{I}_m^{\text{point}}(d_{m,n}) = \begin{cases} 1 & \text{if } d_{m,n} \geq \tau_m \\ 0 & \text{otherwise} \end{cases} \tag{3.7}$$

For area defects (gaps and overlaps), the indicator function evaluates the maximum width w_j of defect polygon j against the threshold:

$$\mathbf{I}_m^{\text{area}}(w_j) = \begin{cases} 1 & \text{if } w_j \geq \tau_m \\ 0 & \text{otherwise} \end{cases} \tag{3.8}$$

The distinction between these formulations reflects the different geometric representations of the defect types. Point defects are evaluated at element centroids and have a single scalar value per element. Area defects are polygonal regions that may intersect multiple elements; the threshold comparison uses the defect's maximum width (a global attribute of the polygon), while the spatial distribution of the penalty is determined by the element intersection calculation described in Section 3.2.4.

3.2.3 SEVERITY FUNCTION

The severity function quantifies the magnitude of threshold exceedance, assigning greater penalty to defects that significantly exceed the threshold compared to those that marginally exceed it. This graduated penalty reflects the physical reality that a gap twice the allowable width is more detrimental to part quality than one barely exceeding the limit.

The severity function takes a quadratic form:

$$s_m(d) = \left(\frac{d}{\tau_m}\right)^2 \tag{3.9}$$

where d is the defect value (either $d_{m,n}$ for point defects or w_j for area defects) and τ_m is the threshold. The severity function is applied only when the defect exceeds the threshold; defects below threshold have zero contribution to the severity penalty.

The quadratic dependence produces rapidly increasing penalties as defects exceed the threshold, as illustrated in Figure 3.4. At the threshold boundary ($d = \tau_m$), the severity equals unity. A defect at twice the threshold incurs four times the penalty, and one at three times the threshold incurs nine times the penalty. This nonlinear scaling ensures that defects far exceeding the threshold dominate the severity penalty, directing attention to the most critical manufacturing violations.

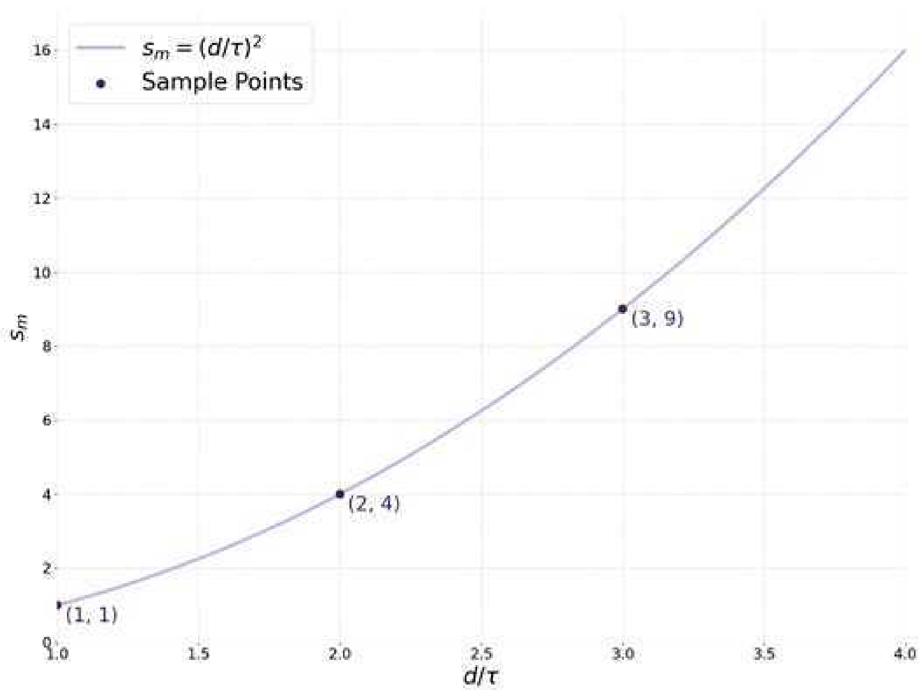

Figure 3.4: Severity function $s_m = (d/\tau_m)^2$ showing quadratic penalty growth for defects exceeding the threshold.

Combined with the frequency penalty (which treats all threshold-exceeding defects equally), the severity penalty provides a balanced assessment that accounts for both the spatial extent and magnitude of defects.

3.2.4 ELEMENT CONTRIBUTION FUNCTIONS

The element contribution function $\delta_{m,n}$ defines how element n contributes to the penalty for defect type m. The formulation differs between point and area defects, reflecting their distinct geometric representations.

3.2.4.1 POINT DEFECT CONTRIBUTION

For point defects, the defect value $d_{m,n}$ is evaluated at the element centroid. If the value exceeds the threshold, the element contributes its full area to the penalty:

$$\delta_{m,n}^{\text{point}} = A_n \cdot \mathbf{I}_m^{\text{point}}(d_{m,n}) \tag{3.10}$$

where A_n is the area of element n and $\mathbf{I}_m^{\text{point}}(d_{m,n})$ is the threshold indicator from Equation (3.7). Elements where the defect value is below threshold contribute zero.

3.2.4.2 AREA DEFECT CONTRIBUTION

For area defects, the contribution depends on the spatial intersection between the defect polygon and the element. Let D_j denote defect polygon j with maximum width w_j, and let $\alpha_{j,n}$ denote the coverage fraction defined in Equation (3.4). The element contribution is:

$$\delta_{m,n}^{\text{area}} = \alpha_{j,n} \cdot A_n \cdot \mathbf{I}_m^{\text{area}}(w_j) \tag{3.11}$$

The product $\alpha_{j,n} \cdot A_n$ equals the absolute intersection area between the defect polygon and the element. This formulation ensures that only the portion of the element actually covered by the defect contributes to the penalty, providing spatial resolution that reflects the true extent of the manufacturing defect.[1]

[1] If multiple defects of the same type were to intersect an element, their contributions would be additive. This situation is physically rare due to the geometric constraints of adjacent courses, where each inter-course region typically contains at most one gap or overlap.

3.2.5 PENALTY FORMULATION

The penalty for each defect type comprises two components: a frequency penalty that measures the fraction of ply area affected by threshold-exceeding defects, and a severity penalty that incorporates the magnitude of threshold exceedance.

3.2.5.1 FREQUENCY PENALTY

The frequency penalty P_m^{freq} is the ratio of total defect-affected area to total ply area:

$$P_m^{\mathrm{freq}} = \frac{1}{A_{\mathrm{ply}}} \sum_{n=1}^{N_e} \delta_{m,n} \tag{3.12}$$

where $\delta_{m,n}$ is the element contribution from Equation (3.10) or (3.11) depending on defect type, and A_{ply} is the total ply area from Equation (3.1). The frequency penalty satisfies $0 \leq P_m^{\mathrm{freq}} \leq 1$, with $P_m^{\mathrm{freq}} = 0$ indicating no defects exceed the threshold and $P_m^{\mathrm{freq}} = 1$ indicating the entire ply area is affected.

3.2.5.2 SEVERITY PENALTY

The severity penalty P_m^{sev} incorporates the severity function to weight contributions by the magnitude of threshold exceedance. The formulation differs between point and area defects, consistent with the element contribution functions defined in Section 3.2.4.

For point defects, the severity function is evaluated at the element-level defect value $d_{m,n}$:

$$P_m^{\mathrm{sev}} = \frac{1}{A_{\mathrm{ply}}} \sum_{n=1}^{N_e} s_m(d_{m,n}) \cdot \delta_{m,n}^{\mathrm{point}} \tag{3.13}$$

For area defects, the severity function is evaluated at the maximum width w_j of defect polygon j, which is a polygon-level attribute applied uniformly across all elements intersected by that polygon:

$$P_m^{\text{sev}} = \frac{1}{A_{\text{ply}}} \sum_{n=1}^{N_e} s_m(w_j) \cdot \delta_{m,n}^{\text{area}} \tag{3.14}$$

Unlike the frequency penalty, P_m^{sev} is not bounded above by 1, as defects significantly exceeding the threshold contribute severity values greater than 1. The implications of this for the overall score bounds are addressed in Section 3.2.7.

3.2.5.3 PHYSICAL INTERPRETATION

The frequency and severity penalties capture complementary aspects of manufacturability:

- **Frequency** answers: *"What fraction of the ply is affected by defects?"*
- **Severity** answers: *"How badly do the defects exceed their thresholds?"*

A ply with many small threshold exceedances will have high frequency but moderate severity. A ply with a single large exceedance will have low frequency but high severity. The weighting scheme allows the user to balance these considerations according to process requirements.

3.2.6 SCORE AGGREGATION

The final MIDS is computed by aggregating the frequency and severity penalties across all defect types, weighted by the user-defined criteria weights.

For each defect type m, the weighted penalty combines the frequency and severity components:

$$P_m = \omega_m^{\text{freq}} \cdot P_m^{\text{freq}} + \omega_m^{\text{sev}} \cdot P_m^{\text{sev}} \tag{3.15}$$

The total manufacturing penalty is then computed as the sum over all defect types in

$\mathcal{M}$:

$$\text{Penalty}_{\text{mfg}} = \sum_{m \in \mathcal{M}} P_m = \sum_{m \in \mathcal{M}} \left(\omega_m^{\text{freq}} \cdot P_m^{\text{freq}} + \omega_m^{\text{sev}} \cdot P_m^{\text{sev}} \right) \tag{3.16}$$

Finally, the MIDS is defined as the clamped complement of the total penalty:

$$\text{MIDS} = \max \left(0,\ 100 - \text{Penalty}_{\text{mfg}} \right) \tag{3.17}$$

This formulation ensures that higher scores correspond to better manufacturability: a score of 100 indicates no penalty (all defects below threshold), while lower scores reflect increasing manufacturing risk. The lower clamp at zero is necessary because the severity penalty is unbounded above, as discussed in Section 3.2.7, and prevents the score from becoming negative in cases of severe process violation.

3.2.7 INTERPRETATION AND BOUNDS

The MIDS formulation guarantees that the score lies within the interval $[0, 100]$, providing an intuitive metric.

3.2.7.1 SCORE BOUNDS

The upper bound is achieved when no defects exceed their thresholds:

$$\delta_{m,n} = 0 \ \forall\, m, n \quad \Rightarrow \quad P_m^{\text{freq}} = P_m^{\text{sev}} = 0 \quad \Rightarrow \quad \text{MIDS} = 100 \tag{3.18}$$

The lower bound of zero is enforced by the clamp in Equation (3.17). While the frequency penalty satisfies $0 \leq P_m^{\text{freq}} \leq 1$ by construction, the severity penalty P_m^{sev} is unbounded above due to the quadratic growth of $s_m(d)$ for defects significantly exceeding the threshold. In cases of severe process violation, $\text{Penalty}_{\text{mfg}}$ can therefore exceed 100, and the clamp ensures the score remains non-negative. A clamped score of zero indicates process conditions with severe and widespread defect violations that necessitate fundamental revision of the process plan. In practice, well-formed process plans produce scores well above zero, and the clamp is rarely active.

3.2.7.2 WEIGHT CUSTOMIZATION

The AHP-based weight assignment allows users to tailor the scoring to specific process priorities. For applications where fiber alignment is critical (e.g., highly loaded primary structure), the angle deviation weight can be increased. For applications with tight surface quality requirements, gap and overlap weights can be prioritized. This flexibility ensures that the MIDS reflects the manufacturing concerns most relevant to the application at hand.

The MIDS provides a foundation for automated process planning by reducing the multidimensional defect space to a single objective function. This enables systematic comparison of ply scenarios and forms the basis for the optimization methodology presented in Section 3.4.

3.3 DESIGN-INFORMED DEFECT SCORE (DIDS)

The MIDS formulation presented in Section 3.2 quantifies manufacturability based solely on defect geometry and threshold exceedance, treating all regions of the ply equally regardless of their structural significance. In practice, however, defects in structurally critical regions pose greater risk to part performance than identical defects in over-designed regions. A gap located where the structure operates near its allowable stress limit is more consequential than the same gap in a region with substantial margin reserve.

The DIDS extends the MIDS framework by incorporating structural performance data from Collier Aerospace's HyperX, an aerospace-grade structural sizing and optimization software [205]. HyperX provides element-level MoS values that quantify the local margin between predicted stress response and allowable limits. By weighting defect penalties according to these margins, the DIDS formulation produces a score that reflects structural criticality in addition to manufacturing quality.

This section presents the data exchange architecture that enables communication between HyperX and CAPP, followed by the mathematical extension of the MIDS penalty formulation to incorporate margin weighting. The resulting DIDS metric provides a structural perspective that complements the manufacturing perspective of MIDS. These two scores are then combined through a weighted formulation to produce the IDS, a unified objective function that supports design-aware process optimization. Application of the complete DIDS and IDS workflow, including data exchange verification and demonstration of the structural feedback loop, is presented in Section 4.3.

3.3.1 DATA EXCHANGE ARCHITECTURE

Communication between HyperX and CAPP is achieved through a standardized, file-based interface that ensures interoperability and unit consistency across both design and manufacturing domains. This interface uses Extensible Markup Language (.xml) files to exchange three primary categories of information: geometric definition, optimization settings, and structural performance data.

The geometric definition is provided by the FEM described in Section 3.1.1. This mesh specifies the structural layout through lists of nodes and elements, where each node entry contains its identification number and spatial coordinates, and each element references the four nodes defining its boundaries along with its centroid position and projected surface area. The mesh serves as the geometric backbone connecting the manufacturing model to the structural analysis environment, enabling one-to-one mapping of elements between HyperX and CAPP.

The optimization settings file contains the process and design parameters required to define the analysis conditions within CAPP. These settings specify allowable layup strategies, defect types, severity thresholds, and weighting coefficients, ensuring that the planning process adheres to the same design rules and constraints enforced during structural analysis.

The MoS data file embeds structural performance information into each element of the mesh. Each record contains element-specific MoS values that quantify the local safety margin between predicted stress response and allowable limits. This data structure converts the mesh into a spatially weighted map of structural sensitivity, allowing CAPP to assess not only the manufacturability of a given layup but also its contribution to overall structural integrity.

All three files are imported directly into CAPP through the software's automated data-exchange module. During the import process, unit alignment, mesh topology, and file compatibility are verified to ensure that the geometry and attribute data from HyperX align precisely with the AFP digital model. Once imported, CAPP visualizes the MoS values across the laminate surface using a carpet plot representation (Figure 3.5). In this visualization, element height indicates the relative margin magnitude, while color gradients—from blue for high-margin regions to red for critical low-margin areas—provide immediate understanding of the structural landscape. This visualization allows process planners to spatially correlate regions of structural importance with the predicted distribution of defects, forming the basis for margin-informed optimization.

3.3.2 FACTOR OF SAFETY TRANSFORMATION

The MoS values provided by HyperX represent the relative margin between applied and allowable stresses, expressed as a ratio. Positive values indicate reserve capacity (the structure can sustain additional load), while negative values indicate that the applied stress exceeds the allowable limit. Because MoS values can be negative in regions where loads exceed allowable limits, they cannot be directly incorporated into the penalty weighting

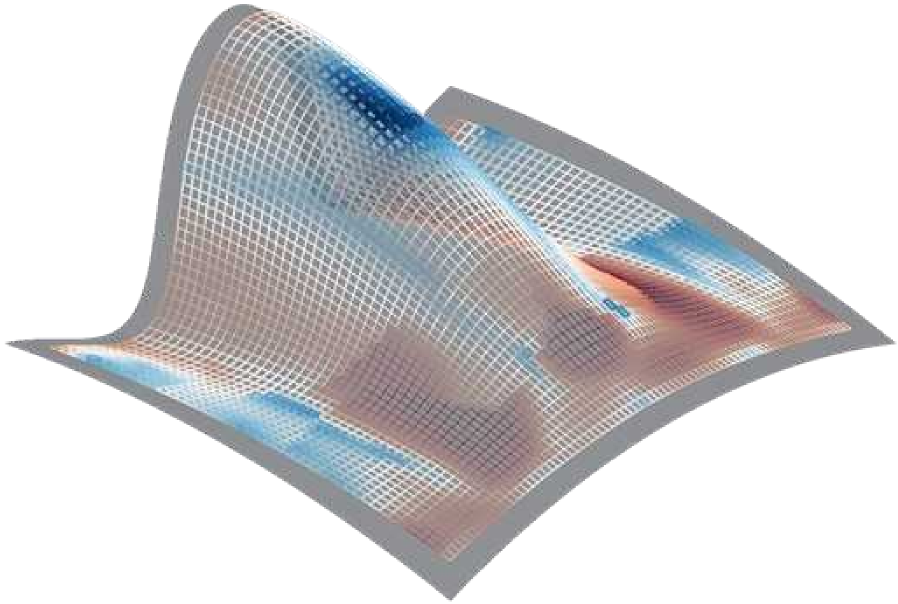

Figure 3.5: Visualization of MoS data within CAPP using carpet plot representation. Element height corresponds to margin magnitude; color indicates structural criticality (blue: high margin, red: low margin).

function without altering the mathematical behavior of the scoring algorithm.

To resolve this, the MoS values are transformed into a positive, nondimensional Factor-of-Safety (FoS) defined as:

$$\mathrm{FoS}_n = 1 + \mathrm{MoS}_n \tag{3.19}$$

This transformation preserves the physical interpretation of margins while maintaining numerical stability. An element with $\mathrm{MoS}_n = 0$ yields $\mathrm{FoS}_n = 1$, indicating that the applied load exactly equals the allowable limit. Higher values correspond to over-designed regions with reserve capacity, while values less than one indicate structurally critical regions where the allowable limit has been exceeded. The formulation assumes $\mathrm{MoS}_n > -1$ for all elements, which corresponds to the physically meaningful operating range; a margin below -1 would indicate a load more than twice the allowable, a condition that necessitates design revision prior to process planning.

3.3.3 MARGIN-WEIGHTED PLY AREA

The margin-weighted penalty formulation requires a normalization basis that reflects the structural significance of each element. The margin-weighted ply area serves this purpose by scaling each element's contribution inversely with its FoS:

$$A_{\text{ply,margin}} = \sum_{n=1}^{N_e} \frac{A_n}{\text{FoS}_n} \tag{3.20}$$

This formulation amplifies the influence of low-margin regions (where FoS_n is small) while attenuating the contribution of structurally robust areas (where FoS_n is large). Physically, this reflects the engineering principle that defects in critical regions warrant greater attention than defects in regions with substantial reserve capacity.

3.3.4 MARGIN-WEIGHTED ELEMENT CONTRIBUTIONS

The element contribution functions defined in Section 3.2.4 are extended to incorporate margin weighting by dividing each contribution by the element's FoS. This modification ensures that defects in structurally critical regions receive proportionally greater penalty than identical defects in over-designed regions.

3.3.4.1 POINT DEFECT CONTRIBUTION

For point defects, the margin-weighted element contribution is:

$$\delta^{\text{point}}_{m,n,\text{margin}} = \frac{A_n}{\text{FoS}_n} \cdot \mathbf{I}^{\text{point}}_m(d_{m,n}) \tag{3.21}$$

where $\mathbf{I}^{\text{point}}_m(d_{m,n})$ is the threshold indicator function from Equation (3.7). Comparing this to the manufacturing contribution in Equation (3.10), the only modification is the division by FoS_n.

3.3.4.2 AREA DEFECT CONTRIBUTION

For area defects, the margin-weighted element contribution is:

$$\delta^{\text{area}}_{m,n,\text{margin}} = \frac{\alpha_{j,n} \cdot A_n}{\text{FoS}_n} \cdot \mathbf{I}^{\text{area}}_m(w_j) \tag{3.22}$$

where $\alpha_{j,n}$ is the coverage fraction from Equation (3.4) and $\mathbf{I}^{\text{area}}_m(w_j)$ is the threshold indicator from Equation (3.8). As with point defects, the margin weighting is introduced through division by FoS_n.

3.3.5 MARGIN-WEIGHTED PENALTY FORMULATION

The margin-weighted penalties follow the same frequency-severity structure established in Section 3.2.5, with the margin-weighted contributions replacing the manufacturing contributions and the margin-weighted ply area replacing the total ply area as the normalization basis.

3.3.5.1 MARGIN-WEIGHTED FREQUENCY PENALTY

The margin-weighted frequency penalty measures the fraction of the margin-weighted ply area affected by threshold-exceeding defects:

$$P^{\text{freq}}_{m,\text{margin}} = \frac{1}{A_{\text{ply,margin}}} \sum_{n=1}^{N_e} \delta_{m,n,\text{margin}} \tag{3.23}$$

where $\delta_{m,n,\text{margin}}$ is the margin-weighted element contribution from Equation (3.21) or (3.22) depending on defect type.

3.3.5.2 MARGIN-WEIGHTED SEVERITY PENALTY

The margin-weighted severity penalty incorporates the severity function to weight contributions by the magnitude of threshold exceedance. As with the MIDS severity penalty, the formulation differs between point and area defects.

For point defects, the severity function is evaluated at the element-level defect value $d_{m,n}$:

$$P_{m,\text{margin}}^{\text{sev}} = \frac{1}{A_{\text{ply,margin}}} \sum_{n=1}^{N_e} s_m(d_{m,n}) \cdot \delta_{m,n,\text{margin}}^{\text{point}} \tag{3.24}$$

For area defects, the severity function is evaluated at the maximum width w_j of defect polygon j, applied uniformly across all elements intersected by that polygon:

$$P_{m,\text{margin}}^{\text{sev}} = \frac{1}{A_{\text{ply,margin}}} \sum_{n=1}^{N_e} s_m(w_j) \cdot \delta_{m,n,\text{margin}}^{\text{area}} \tag{3.25}$$

The threshold indicator function, severity function, and defect criteria weights remain unchanged from the MIDS formulation; only the element contributions and normalization basis differ.

3.3.6 DESIGN-INFORMED DEFECT SCORE

The margin-weighted penalties are aggregated following the same procedure used for MIDS in Section 3.2.6. For each defect type m, the weighted margin penalty is:

$$P_{m,\text{margin}} = \omega_m^{\text{freq}} \cdot P_{m,\text{margin}}^{\text{freq}} + \omega_m^{\text{sev}} \cdot P_{m,\text{margin}}^{\text{sev}} \tag{3.26}$$

The total margin penalty is the sum over all defect types:

$$\text{Penalty}_{\text{margin}} = \sum_{m \in \mathcal{M}} P_{m,\text{margin}} \tag{3.27}$$

The DIDS is then defined as the clamped complement of the total margin penalty:

$$\text{DIDS} = \max\left(0,\ 100 - \text{Penalty}_{\text{margin}}\right) \tag{3.28}$$

The lower clamp at zero is necessary for the same reason as in the MIDS formulation: the quadratic severity function is unbounded above, and the clamp ensures the score remains non-negative in cases of severe process violation. The DIDS is otherwise analogous to MIDS but reflects structural criticality rather than pure manufacturing quality. A DIDS of 100 indicates no threshold-exceeding defects in any structurally significant region, while lower values indicate increasing penalty contributions weighted by structural importance.

3.3.7 INTEGRATED DEFECT SCORE

The IDS combines the manufacturing perspective (MIDS) and the structural perspective (DIDS) through a weighted combination:

$$\text{IDS} = w_{\text{design}} \cdot \text{DIDS} + w_{\text{mfg}} \cdot \text{MIDS} \tag{3.29}$$

where $w_{\text{mfg}}, w_{\text{design}} \geq 0$ are user-defined weighting factors representing the relative importance of manufacturability and structural sensitivity, respectively. These weighting factors satisfy the normalization constraint:

$$w_{\text{mfg}} + w_{\text{design}} = 1 \tag{3.30}$$

This constraint ensures that IDS remains bounded in the interval $[0, 100]$, maintaining consistency with the MIDS and DIDS interpretations.[2]

By adjusting the weighting factors, process planners can prioritize manufacturing

[2] In the CAPP software implementation, the weighting factors are specified as percentages summing to 100 rather than fractions summing to unity.

efficiency ($w_{\text{mfg}} \to 1$), structural safety ($w_{\text{design}} \to 1$), or an intermediate balance. This flexibility transforms CAPP from a defect-avoidance tool into a design-aware optimization platform capable of supporting trade-off analyses between competing objectives.

3.3.8 INTERPRETATION AND TRADE-OFF ANALYSIS

The IDS formulation establishes a unified metric that couples manufacturing feasibility with structural performance. An IDS of 100 corresponds to a defect-free, structurally robust configuration, whereas lower values indicate increasing defect concentration or reduced safety margin in critical regions.

The weighting factors enable explicit trade-off analysis between competing objectives. Setting $w_{\text{mfg}} = 1$ recovers the pure MIDS, treating all regions equally regardless of structural significance. Setting $w_{\text{design}} = 1$ recovers the pure DIDS, producing a score driven entirely by structural criticality that potentially accepts defects in over-designed regions while strictly penalizing defects near allowable limits. Intermediate weight configurations balance these perspectives according to application requirements.

This formulation supports the closed-loop workflow envisioned in the iSPP framework: CAPP optimizes process parameters to maximize IDS, HyperX validates the resulting configuration through structural analysis, and any residual issues are addressed through design modifications or local reinforcement. The IDS metric thus serves as the interface between manufacturing optimization and structural verification, enabling coordinated decision-making across traditionally separate disciplines. Quantitative demonstration of this workflow on the Quartic benchmark surface, including data exchange verification and the complete structural feedback loop from process optimization through reinforcement and reanalysis, is presented in Section 4.3.

3.4 PLY-LEVEL OPTIMIZATION (PLO)

The scoring formulations presented in the preceding sections reduce the multidimensional defect space to a single scalar objective, the IDS, enabling quantitative comparison of candidate ply scenarios. With this metric established, the next challenge is to systematically search the process planning design space to identify the ply scenario that maximizes IDS for a given ply. This search constitutes PLO, the subject of this section.

The central difficulty of PLO is that each candidate evaluation requires a full VCP simulation—path generation, defect prediction, and data export—followed by CAPP scoring. This computational cost, typically on the order of minutes per scenario depending on part size and available hardware, precludes exhaustive enumeration of the continuous design space and necessitates a structured search strategy that balances exploration with computational budget.

This section formalizes the PLO problem, defines the design space and constraints, and presents the two-phase search algorithm implemented within CAPP. The algorithm combines a coarse global survey with a local refinement procedure termed the perpendicular-shift method. The development of the local refinement approach is presented in its full evolution: the original formulation assumed a smooth objective landscape amenable to classical one-dimensional optimization, an assumption that empirical investigation revealed to be invalid. The findings that led to this conclusion are documented, and the resulting sampling-based refinement strategy is presented as the methodologically sound replacement. A spatial diversity criterion for seed selection is introduced to improve robustness against local optima. Application of the PLO algorithm to the Quartic benchmark surface, including baseline characterization, survey results, and optimization outcomes, is presented in Section 4.4.

3.4.1 PROBLEM FORMULATION

3.4.1.1 DESIGN VARIABLES

A ply scenario $\mathbf{x}$ is defined by two decision variables: a starting point $\mathbf{p}_0$ on the tool surface and a layup strategy ℓ. Formally,

$$\mathbf{x} = (\mathbf{p}_0, \ell) \tag{3.31}$$

where $\mathbf{p}_0 \in \mathcal{S}$ is a point on the tool surface $\mathcal{S} \subset \mathbb{R}^3$ and $\ell \in \mathcal{L} = \{\ell_1, \ell_2, \ldots, \ell_{N_\ell}\}$ is a member of the finite set of user-selected layup strategies. The starting point is a continuous variable constrained to the surface, while the layup strategy is a categorical variable drawn from a discrete set.

The starting point coordinates are expressed in the Cartesian frame $\mathbf{p}_0 = (x_0, y_0, z_0)$, with the constraint that $\mathbf{p}_0$ lies on the tool surface. Equivalently, through the surface parameterization, the point can be represented by its parameter-space coordinates (u_0, v_0) such that $\mathbf{p}_0 = \mathbf{r}(u_0, v_0)$, where $\mathbf{r} : \mathbb{R}^2 \to \mathbb{R}^3$ is the surface map. This parameterization reduces the continuous design space from three constrained Cartesian coordinates to two unconstrained surface parameters.

The number of tows per course is a configurable parameter in VCP and can be varied dynamically by most modern AFP machines. However, it is typically selected by the process planner during initial project setup based on the part geometry, material system, and production requirements, and is held constant across ply scenarios for a given optimization study. It is therefore treated as a fixed parameter rather than a decision variable within the PLO formulation.

3.4.1.2 OBJECTIVE FUNCTION

The objective function for PLO is the IDS defined in Equation (3.29):

$$f(\mathbf{x}) = \mathrm{IDS}(\mathbf{p}_0, \ell) = w_{\mathrm{design}} \cdot \mathrm{DIDS}(\mathbf{p}_0, \ell) + w_{\mathrm{mfg}} \cdot \mathrm{MIDS}(\mathbf{p}_0, \ell) \tag{3.32}$$

where w_{design} and w_{mfg} are the user-defined weighting factors satisfying $w_{\mathrm{design}} + w_{\mathrm{mfg}} = 1$. When structural margin data from HyperX is unavailable, $w_{\mathrm{design}} = 0$ and the IDS reduces identically to MIDS, ensuring that the optimization framework remains functional without requiring a structural model.

The objective function is a black-box function: no closed-form expression relates $\mathbf{x}$ to $f(\mathbf{x})$. Each evaluation requires executing the full VCP simulation pipeline to generate courses, predict defects, and map defect geometries onto the finite element mesh, followed by CAPP's scoring computation. The function is therefore computationally expensive, and its analytical properties (continuity, differentiability, convexity) are unknown *a priori*.

3.4.1.3 CONSTRAINTS

The starting point must satisfy two constraints:

1. **Surface constraint:** $\mathbf{p}_0 \in \mathcal{S}$, i.e., the point must lie on the tool surface. This is enforced through surface projection using the parameterization $\mathbf{p}_0 = \mathbf{r}(u_0, v_0)$.
2. **Boundary constraint:** $(u_0, v_0) \in \mathcal{B}$, where $\mathcal{B} \subset \mathbb{R}^2$ is the ply boundary polygon in parameter space. This ensures that the starting point falls within the geometric extent of the ply being optimized.

The boundary constraint is enforced through a containment test: candidate points whose parameter-space projections fall outside $\mathcal{B}$ are discarded. The surface constraint is enforced through projection of Cartesian points onto $\mathcal{S}$ using standard computational geometry routines.

3.4.1.4 OPTIMIZATION STATEMENT

The PLO problem is stated as:

$$\mathbf{x}^* = \arg\max_{\mathbf{x}=(\mathbf{p}_0,\ell)} f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{p}_0 \in \mathcal{S},\ (u_0, v_0) \in \mathcal{B},\ \ell \in \mathcal{L} \tag{3.33}$$

This is a mixed-variable optimization problem: continuous in $\mathbf{p}_0$ and categorical in ℓ. The mixed nature, combined with the black-box objective and per-evaluation computational cost, places PLO in the class of expensive black-box optimization problems with mixed decision variables [206]. The algorithm developed in the following subsections addresses this problem through a structured sampling strategy that separates the categorical variable (handled by enumeration) from the continuous variable (handled by a two-phase global-local search).

3.4.2 ALGORITHM OVERVIEW

The PLO algorithm decomposes the mixed-variable problem into two components. The categorical variable ℓ is handled by explicit enumeration: the survey evaluates all user-selected layup strategies at each candidate starting point. The continuous variable $\mathbf{p}_0$ is searched through a two-phase procedure consisting of a global survey followed by local refinement via perpendicular shifts. Algorithm 1 summarizes the complete procedure.

The following subsections detail each component of the algorithm.

3.4.3 PHASE 1: GLOBAL SURVEY

The global survey provides coarse coverage of the ply domain by evaluating IDS at a regular grid of starting points, each paired with every user-selected layup strategy. This phase serves two purposes: it identifies promising regions of the design space for subsequent

Algorithm 1 Ply-Level Optimization (PLO)

Require: Ply boundary $\mathcal{B}$, tool surface $\mathcal{S}$, layup strategies $\mathcal{L}$, survey spacing Δ_s, number of shift points n, number of seeds K, minimum seed separation $d_{\min}$

Ensure: Optimal ply scenario $\mathbf{x}^*$

Phase 1: Global Survey

Generate grid $\mathcal{G}$ over $\mathcal{B}$ with spacing Δ_s

$C \leftarrow \{(\mathbf{p}_i, \ell_j) : \mathbf{p}_i \in \mathcal{G},\ \ell_j \in \mathcal{L}\}$ ▷ All grid-strategy combinations

Evaluate $f(\mathbf{x})$ for all $\mathbf{x} \in C$ via VCP and CAPP

Rank C by descending IDS

Select K seed scenarios $\mathcal{K}$ using spatial diversity criterion (Section 3.4.4)

Phase 2: Local Refinement (Perpendicular Shift)

for each seed $\mathbf{x}_k \in \mathcal{K}$ **do**
    Generate n perpendicular-shift candidates from $\mathbf{x}_k$ (Section 3.4.5)
    Evaluate all candidates via VCP and CAPP
    while best candidate is at shift range boundary **and** shift range has not reached $\mathcal{B}$ **do**
        Extend shift range in the direction of the boundary optimum
        Generate and evaluate additional candidates
    end while
end for

$\mathbf{x}^* \leftarrow \arg\max_{\mathbf{x}} f(\mathbf{x})$ over all evaluated scenarios
return $\mathbf{x}^*$

local refinement, and it generates a scored population of ply scenarios that can be used as candidates for LLO (Section 3.5).

3.4.3.1 GRID GENERATION

The survey grid is constructed in the surface parameter space (u, v). Let $\mathcal{B}$ denote the ply boundary polygon in parameter space, with bounding box $[u_{\text{min}}, u_{\text{max}}] \times [v_{\text{min}}, v_{\text{max}}]$. The grid spacing Δ_s is defined as a user-specified percentage of the course width:

$$\Delta_s = \frac{\gamma}{100} \cdot W_c \cdot \lambda \tag{3.34}$$

where γ is the spacing parameter (expressed as a percentage, default 150%), $W_c = N_t \cdot w_t$ is the course width (the product of the number of tows N_t and tow width w_t), and λ is a scaling factor that accounts for the distortion between parameter space and physical space:

$$\lambda = \sqrt{\frac{A_{\mathcal{B},\text{param}}}{A_{\mathcal{B},\text{phys}}}} \tag{3.35}$$

Here, $A_{\mathcal{B},\text{param}}$ is the area of the ply boundary polygon in parameter space and $A_{\mathcal{B},\text{phys}}$ is the physical ply area. This scaling ensures that the grid spacing in parameter space corresponds approximately to the intended physical spacing on the tool surface, despite the nonlinear relationship between the two coordinate systems. The factor λ carries units of parameter coordinates per unit length, so that Δ_s is expressed in parameter-space coordinates consistent with the grid coordinate system defined above.

If a baseline ply scenario (scenario 0) exists, the grid is phase-aligned to its parameter-space coordinates $(u_{\text{ref}}, v_{\text{ref}})$. This alignment ensures that the baseline starting point coincides with a grid node and is included in the survey ranking without requiring a redundant VCP evaluation.

Grid nodes are enumerated over the bounding box, and each candidate is subjected to the boundary containment test $(u_i, v_i) \in \mathcal{B}$. Points outside the ply boundary are discarded. The surviving points are projected onto the tool surface to obtain their Cartesian coordinates $\mathbf{p}_i = \mathbf{r}(u_i, v_i)$.

3.4.3.2 SCENARIO ENUMERATION AND EVALUATION

Each retained grid point $\mathbf{p}_i$ is combined with every selected layup strategy $\ell_j \in \mathcal{L}$ to form a candidate ply scenario $\mathbf{x}_{ij} = (\mathbf{p}_i, \ell_j)$. The total number of survey scenarios is:

$$N_{\text{survey}} = N_g \cdot N_\ell \tag{3.36}$$

where N_g is the number of grid points within $\mathcal{B}$ and $N_\ell = |\mathcal{L}|$ is the number of selected layup strategies.

All survey scenarios are submitted to VCP for path generation and defect prediction, using the subprocess integration described in Section 3.1.6. Upon completion, CAPP imports the defect data, computes IDS for each scenario, and ranks the full set by descending score.

3.4.4 SEED SELECTION WITH SPATIAL DIVERSITY

The transition from global survey to local refinement requires selecting a set of seed scenarios that will serve as the starting points for perpendicular-shift searches. Selecting only the single highest-scoring survey scenario risks committing the local search to a single region of the design space that may contain a local, rather than global, optimum. Conversely, selecting the top K scenarios by raw score provides no guarantee of spatial diversity; several high-scoring scenarios may cluster in the same neighborhood and yield redundant local searches.

To balance exploitation of high-scoring regions with exploration of the design space, the seed selection employs a spatial diversity criterion. The procedure, detailed in Algorithm 2, selects seeds iteratively: the highest-scoring scenario is always chosen first, and each subsequent seed must satisfy a minimum separation distance from all previously selected seeds in parameter space.

Algorithm 2 Seed selection with spatial diversity criterion

Require: Ranked survey scenarios $\mathcal{R} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{N_{\text{survey}}}\}$ (descending IDS), number of seeds K, minimum separation $d_{\min}$
Ensure: Seed set $\mathcal{K}$ with $|\mathcal{K}| \leq K$
$\mathcal{K} \leftarrow \{\mathbf{x}_1\}$ ▷ Best-scoring scenario is always a seed
for $i = 2$ to N_{survey} **do**
  if $|\mathcal{K}| = K$ **then**
    break
  end if
  $d_i \leftarrow \min_{\mathbf{x}_k \in \mathcal{K}} \|(u_i, v_i) - (u_k, v_k)\|_2$
  if $d_i \geq d_{\min}$ **then**
    $\mathcal{K} \leftarrow \mathcal{K} \cup \{\mathbf{x}_i\}$
  end if
end for
return $\mathcal{K}$

The minimum separation distance $d_{\min}$ is defined as a fraction of the survey grid spacing:

$$d_{\min} = \beta \cdot \Delta_s \tag{3.37}$$

where $\beta \in (0, 1]$ is a user-configurable parameter (default $\beta = 0.5$). This choice ensures that seeds are at least half a grid cell apart in parameter space, preventing trivially adjacent seeds while permitting multiple seeds within the same broad region if that region contains several distinct high-scoring zones.

The spatial diversity criterion is motivated by the non-smooth nature of the IDS landscape, documented in Section 3.4.6. Because nearby starting points may share the same local basin of attraction with respect to defect formation, requiring spatial separation increases the probability that the seed set covers distinct basins and that at least one local refinement converges to a neighborhood containing the global optimum.

If the ranked list is exhausted before K seeds are selected (i.e., all remaining scenarios violate the separation criterion), the algorithm terminates with $|\mathcal{K}| < K$ seeds. This can occur when the ply boundary is small relative to the grid spacing, resulting in fewer viable candidate regions.

3.4.5 PHASE 2: PERPENDICULAR-SHIFT REFINEMENT

Once the seed scenarios are identified, each seed undergoes local refinement through the perpendicular-shift method. This method exploits a key property of AFP path generation: translating the starting point along the course direction does not alter the underlying reference curve geometry and therefore produces identical defect distributions. Meaningful variation in defect patterns arises only when the starting point is displaced perpendicular to the course direction, which shifts the entire course pattern laterally relative to the ply boundary and modifies the locations of course-to-course gap and overlap formations. The perpendicular-shift method therefore restricts the local search to the single degree of freedom that influences the objective function.

3.4.5.1 PERPENDICULAR DIRECTION COMPUTATION

For a seed scenario $\mathbf{x}_k = (\mathbf{p}_k, \ell_k)$, the perpendicular direction is computed on the tool surface at the starting point $\mathbf{p}_k$. Let $\hat{\mathbf{n}}(\mathbf{p}_k)$ denote the outward surface normal at $\mathbf{p}_k$, and let θ denote the nominal ply fiber angle. The tangent direction of the reference curve at $\mathbf{p}_k$ is determined by θ and the surface normal, and the perpendicular direction $\hat{\mathbf{b}}$ is the binormal vector lying in the tangent plane of the surface, perpendicular to the reference curve tangent:

$$\hat{\mathbf{b}} = \mathbf{R}_{\text{rosette}} \cdot \text{binormal}(\hat{\mathbf{n}}(\mathbf{p}_k),\ \theta) \tag{3.38}$$

where $\mathbf{R}_{\text{rosette}}$ is the rotation matrix associated with the laminate rosette coordinate system. The rosette rotation accounts for any angular offset between the global coordinate system and the laminate reference frame, ensuring that the perpendicular direction is consistent with the fiber angle convention used by VCP.

3.4.5.2 CANDIDATE GENERATION

The perpendicular-shift candidates are generated by translating the seed starting point along $\hat{\mathbf{b}}$ by prescribed distances. The shift range spans one course width, centered on the seed point. For n candidate points (where n is odd to include the seed and symmetric pairs about it), the shift distances are:

$$d_i = \left(\frac{i - (n-1)/2}{n-1}\right) \cdot W_c, \quad i = 0, 1, \ldots, n-1 \tag{3.39}$$

where W_c is the course width defined previously. This produces n evenly spaced values ranging from $-W_c/2$ to $+W_c/2$, with $d_{(n-1)/2} = 0$ corresponding to the seed point itself. Each shift distance generates a candidate point:

$$\mathbf{p}_i = \mathbf{p}_k + d_i \cdot \hat{\mathbf{b}}, \quad i = 0, 1, \ldots, n-1 \tag{3.40}$$

The resulting candidates are projected onto the tool surface to ensure they satisfy the surface constraint. The layup strategy is inherited from the seed scenario and held constant during the perpendicular-shift phase. The candidate at $i = (n-1)/2$ corresponds to the seed point and does not require a new VCP evaluation, as its score is already available from the survey phase. In total, $n-1$ new VCP evaluations are required per seed.

3.4.5.3 ITERATIVE EDGE EXTENSION

After scoring all perpendicular-shift candidates, the algorithm examines whether the highest-scoring candidate lies at an extremity of the shift range. If the optimum is interior to the sampled set, the local refinement is considered complete for that seed. If the optimum lies at the boundary of the shift range, the search may be prematurely truncated, and additional sampling is warranted.

In this case, the shift range is extended in the direction of the boundary optimum. A new set of candidates is generated, centered on the extremal point, spanning an additional course width in the direction indicated. This extension-evaluation cycle repeats until one of two termination conditions is met:

1. The highest-scoring candidate is interior to the current shift range.
2. The shift range reaches the ply boundary $\mathcal{B}$, and no further extension is geometrically possible.

The iterative extension ensures that the local refinement is not artificially constrained by the initial shift range while naturally terminating at the ply boundary if the score continues to improve in one direction.

3.4.6 SMOOTHNESS ANALYSIS AND METHODOLOGICAL IMPLICATIONS

The perpendicular-shift method was originally conceived as a one-dimensional optimization along the transverse direction, with the expectation that the IDS would vary smoothly as a function of shift distance d. Under this assumption, the discrete samples from Equation (3.40) would be used to construct an interpolating function $\hat{f}(d)$, and a classical root-finding method such as Brent's method [207] could locate the precise maximum without requiring exhaustive sampling.

Brent's method combines inverse quadratic interpolation, secant method, and bisection to find the root of a function (or equivalently, the extremum of a unimodal function) within a bracketed interval. Given a sufficiently smooth objective and three initial points that bracket the maximum, the method achieves superlinear convergence without requiring derivative information, making it well suited for expensive black-box functions where gradient evaluations are unavailable [208].

To evaluate this approach, the perpendicular-shift phase was executed with a dense sampling of candidates, and the resulting IDS values were plotted against shift distance for representative plies. Figure 3.6 presents these profiles.

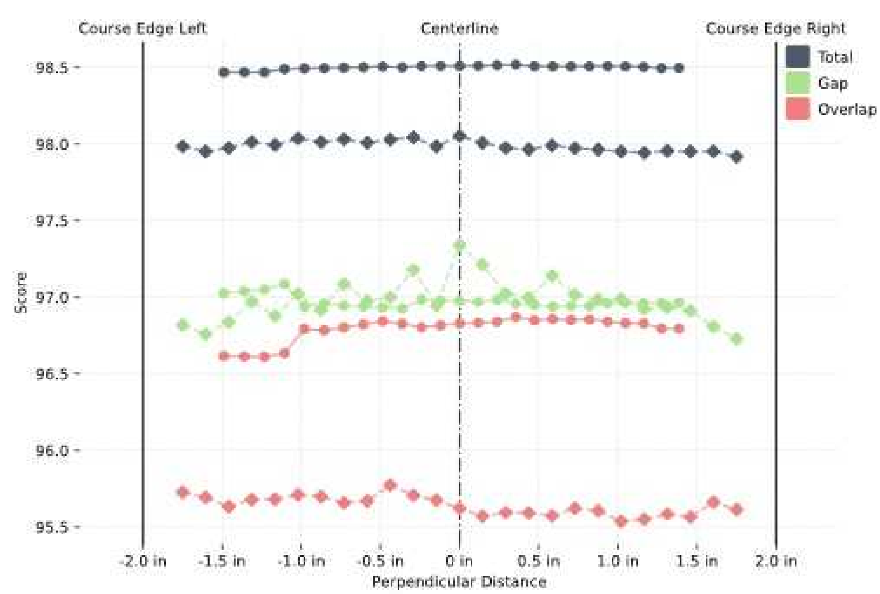

Figure 3.6: IDS component scores versus perpendicular shift distance. Gray lines indicate total scores; green and red lines correspond to gap and overlap component scores, respectively. Dots represent a 90° ply; diamonds represent a 45° ply. Score discontinuities violate the smoothness assumption required for interpolation-based optimization.

The results reveal that the IDS does not vary smoothly as a function of perpendicular shift distance. The score profiles exhibit discontinuous jumps, particularly in the gap and overlap penalty components. These discontinuities are attributed to the discrete nature of tow add and drop events at course boundaries: as the starting point shifts perpendicular to the course direction, individual tows at the edges of the course periodically enter or exit the ply boundary. Each such event introduces or removes an entire tow's width of material coverage, producing a step change in the gap and overlap defect distribution and a corresponding discontinuity in the score.

This finding has a direct methodological consequence. Because the objective function is discontinuous with respect to the perpendicular shift parameter, interpolation-based optimization methods—including Brent's method, polynomial curve fitting, and gradient-based approaches—cannot be reliably applied. An interpolant constructed through discontinuous data will produce spurious extrema that do not correspond to actual optima of the true objective. The application of such methods would therefore introduce systematic error rather than improving precision.

Based on this analysis, the PLO algorithm adopts a sampling-based approach for the local refinement phase: the perpendicular shift generates a discrete set of candidates, evaluates each through VCP and CAPP, and selects the candidate with the highest IDS. This approach makes no assumptions about the smoothness or continuity of the objective function and is therefore robust to the discontinuities observed in the empirical data. The number of perpendicular-shift candidates n controls the resolution of the local search and can be increased by the user at the expense of additional computation time.

3.4.7 COMPUTATIONAL CONSIDERATIONS

The total number of VCP evaluations required by the PLO algorithm is:

$$N_{\text{eval}} = N_{\text{survey}} + K \cdot (n - 1 + N_{\text{ext}}) \tag{3.41}$$

where $N_{\text{survey}} = N_g \cdot N_\ell$ is the number of survey evaluations from Equation (3.36), K is the number of seeds, $n - 1$ is the number of new perpendicular-shift candidates per seed (excluding the already-evaluated seed), and N_{ext} is the number of additional evaluations from edge extension iterations (zero if no extensions are triggered). The computational cost scales linearly with each of these parameters and with the per-evaluation cost of VCP, which depends on part size, mesh density, and course width.

The survey spacing parameter γ provides the primary mechanism for controlling computational effort. Larger values of γ produce coarser grids with fewer survey points, reducing total computation time at the risk of missing localized optima. Smaller values increase grid density and improve coverage at the cost of additional evaluations. In practice, the default spacing provides a reasonable trade-off between coverage and computation time for the geometries encountered in this work. Table 3.4 summarizes the user-configurable parameters and their default values.

Table 3.4: Summary of PLO algorithm parameters.

Parameter	**Symbol**	**Default**	**Description**
Survey spacing	γ	150%	Grid spacing (% of course width)
Shift points	n	7	Candidates per perpendicular shift
Number of seeds	K	5	Seeds for local refinement
Seed separation	β	0.5	Min. separation (fraction of Δ_s)
Layup strategies	$\mathcal{L}$	User-selected	Strategies evaluated per grid point

3.4.8 DUAL ROLE OF PLO

The PLO algorithm serves two complementary purposes within the iSPP framework. In its primary role, PLO optimizes an individual ply by searching for the starting point and layup strategy combination that maximizes IDS. This capability is valuable for plies with highly non-uniform boundaries or complex surface geometry, where the choice of starting point can meaningfully influence defect distribution.

In its secondary role, PLO functions as a scenario generation engine for LLO (Section 3.5). The global survey and local refinement phases produce a diverse population of scored ply scenarios, each representing a distinct defect distribution. These scenarios constitute the candidate pool from which LLO selects optimal ply combinations to minimize through-thickness defect stacking. The quality and diversity of this candidate pool directly influences the effectiveness of LLO; a well-distributed set of scenarios provides LLO with more options for achieving favorable defect staggering across the laminate.

This dual role reflects an important insight that emerged from the development of the PLO methodology: for a single ply evaluated in isolation, the defect distribution is primarily governed by the tool surface geometry, and starting point variation has limited ability to alter the fundamental defect pattern. However, when the same starting-point variation is considered across an entire laminate, it provides the mechanism for defect staggering—ensuring that defects in adjacent plies do not align through the thickness. The value of PLO therefore increases substantially when coupled with LLO, as the scenario diversity generated by PLO provides LLO with a richer candidate pool from which to construct staggering configurations across the laminate thickness, as demonstrated in Section 4.5.

3.5 LAMINATE-LEVEL OPTIMIZATION (LLO)

The PLO methodology presented in Section 3.4 optimizes individual plies by searching for the starting point and layup strategy combination that maximizes the IDS. However, PLO evaluates each ply in isolation and does not account for how its defect distribution aligns with those of other plies in the laminate. Experimental studies have demonstrated that defect configurations exert a substantially greater effect on mechanical performance at the laminate level—with strength reductions of up to 15%—than at the ply level, where reductions remain below 5% [136]. Among the various stacking configurations, defects aligned directly through the thickness constitute the most structurally detrimental case: they produce severe out-of-plane fiber waviness, concentrated resin-rich regions, and localized thickness variations that serve as primary sites for failure initiation and delamination propagation [140]. LLO addresses this laminate-level concern by optimizing the distribution and dispersion of defects across the full stacking sequence, with the specific objective of minimizing through-thickness stacking concentrations.

Industrial practice addresses this objective through ad hoc stagger shifts, in which a process planner manually displaces starting points of similarly oriented plies by a fraction of the course width [52]. The LLO methodology presented in this section replaces this manual practice with a systematic, data-driven algorithm that quantifies stacking through the element-level scoring infrastructure established in Sections 3.2–3.3 and searches the combinatorial space of ply scenario assignments using a coordinate-wise optimization strategy.

Two complementary severity kernels are developed within a unified LLO framework. The first, designated the *stagger shift kernel*, models through-thickness defect interaction as a proximity-weighted coincident stacking measure: it penalizes defects only when they land at the same element across multiple plies, with the penalty magnitude modulated by the ply separation. The second, designated the *gravity-based kernel*, derives from an analogy to Newton's law of universal gravitation: it treats each defective element as a point mass whose defect density determines its gravitational attraction to other defective elements, producing a spatially-decaying interaction that penalizes both direct stacking and near-miss proximity. The two kernels capture complementary aspects of defect arrangement. The coincident stacking kernel targets the direct through-thickness alignment that is of greatest structural concern, while the gravity-based kernel additionally captures the spatial concentration of defects across a local neighborhood. Both kernels operate within the same L-IDS scoring pipeline, are optimized using the same cyclic coordinate algorithm, and produce scores on the same scale bounded above by 100, enabling direct comparison of the two methods on any given laminate configuration.

Section 3.5.1 defines the design variables, candidate pool construction modes, and combinatorial complexity of the problem. Section 3.5.2 extends the IDS framework to the laminate level through element-level defect accumulation and the two severity kernels, followed by the penalty formulation and score aggregation procedure. Section 3.5.3 presents the cyclic coordinate optimization algorithm and establishes its convergence properties. Section 3.5.4 provides the rationale for this algorithmic choice relative to metaheuristic alternatives. Section 3.5.5 characterizes the scope of defect interactions captured by the two kernels. Section 3.5.6 analyzes the computational cost of each phase. Application of the LLO algorithm to the Quartic benchmark surface, including baseline and optimized laminate configurations, defect stacking comparisons, and a head-to-head evaluation of the two severity kernels, is presented in Section 4.5.

3.5.1 PROBLEM FORMULATION

3.5.1.1 DESIGN VARIABLES

A laminate scenario $\mathbf{L}$ is defined by a vector of ply scenario assignments:

$$\mathbf{L} = (s_1, s_2, \ldots, s_{N_p}) \tag{3.42}$$

where N_p is the number of plies in the laminate and each $s_i \in \mathcal{P}_i$ is the ply scenario assigned to ply i. A ply scenario encapsulates a starting point and layup strategy pair, as defined in Equation (3.31), together with its associated defect distribution pre-computed by VCP. At the LLO level, each ply scenario is treated as an atomic unit selected from a finite candidate pool; the continuous starting-point search has already been performed during candidate generation.

3.5.1.2 CANDIDATE POOL CONSTRUCTION

The candidate pool $\mathcal{P}_i$ for ply i is constructed through one of two mechanisms, depending on the available pre-existing data.

The first method, Mode A, utilizes a pre-existing laminate scenario when available, for example from a prior PLO run or from manual process planning. LLO generates perpendicular-shift candidates centered on the existing starting point for each ply. The shift procedure is identical to the one described in Section 3.4.5: candidates are generated by translating the starting point along the binormal direction in increments spanning one course width. The number of shift candidates defaults to $n_i = N_{\text{tows},i} - 1 + (N_{\text{tows},i} \bmod 2)$, which yields the largest odd integer less than or equal to $N_{\text{tows},i}$, ensuring the candidate set includes the seed point and symmetric shift pairs spanning one course width. This default is user-configurable, allowing increased resolution at the cost of additional VCP evaluations. Each shifted candidate is subject to a boundary containment check; candidates whose projected parameter-space coordinates fall outside the ply boundary $\mathcal{B}_i$ are discarded. The pool $\mathcal{P}_i$ then consists of the original baseline scenario plus all valid shifted variants.

The second method, Mode B, is used when starting from scratch and creates a PLO-derived pool of candidate scenarios. When no laminate scenario exists, PLO is automatically performed, and the candidate pool consists of all scored ply scenarios generated during the PLO survey and perpendicular-shift phases (Section 3.4.3–3.4.5). This mode provides a spatially diverse set of candidates drawn from the full extent of the PLO search, including scenarios at multiple grid locations and across multiple layup strategies.

In both modes, the baseline scenario—corresponding to the starting point used for structural margin calculations in HyperX—is always included in $\mathcal{P}_i$. This guarantees that the optimization cannot produce a result worse than the baseline configuration. The pool construction is independent of the severity kernel selected for the optimization, so a single pool generation serves both the stagger shift and gravity-based variants.

3.5.1.3 OBJECTIVE FUNCTION

The objective is to maximize the L-IDS:

$$\mathbf{L}^* = \arg\max_{\mathbf{L}} \ \text{L-IDS}(\mathbf{L}) \tag{3.43}$$

where L-IDS is the laminate-level integrated defect score defined in Section 3.5.2. This function extends the IDS framework to the laminate level by aggregating through-thickness defect stacking penalties across all plies.

As with the ply-level objective, L-IDS is a black-box function: no closed-form expression relates the ply scenario assignments to the score. Each evaluation requires element-level defect accumulation across all plies and computation of the stacking penalties described below. However, because L-IDS operates on pre-computed defect data rather than requiring new VCP simulations, individual evaluations are computationally inexpensive (on the order of seconds).

3.5.1.4 COMBINATORIAL COMPLEXITY

The exhaustive search space has cardinality

$$|\Omega| = \prod_{i=1}^{N_p} |\mathcal{P}_i| \tag{3.44}$$

For laminates with even modest numbers of plies and candidates per ply, this space grows rapidly. As a concrete illustration, an 11-ply laminate with 7 candidates per ply yields $|\Omega| = 7^{11} \approx 1.98 \times 10^9$ possible laminate scenarios. Exhaustive enumeration is therefore intractable, necessitating a structured search strategy that exploits the decomposable nature of the problem.

3.5.2 LAMINATE-LEVEL SCORING

The laminate-level score extends the ply-level IDS framework to detect and penalize through-thickness defect stacking. Following the naming convention established for the ply-level metrics, the laminate-level scores are designated Laminate-level Manufacturing-Informed Defect Score (L-MIDS), Laminate-level Design-Informed Defect Score (L-DIDS), and L-IDS for the manufacturing, design-informed, and integrated variants, respectively. These metrics are architecturally parallel to their ply-level counterparts—employing the same frequency-severity penalty structure and AHP-derived defect weights—but operate on accumulated through-thickness defect data through one of two severity kernels that quantify the spatial concentration of defects across the stacking sequence. The scoring operates on the shared finite element mesh (Section 3.1.1), leveraging the element-level defect data already computed for each ply scenario during PLO or candidate generation.

The scoring formulation has a modular structure. The element-level defect accumulation (Section 3.5.2.1) and the frequency penalty are shared across both severity kernels. The two kernels—the stagger shift coincident stacking measure (Section 3.5.2.2) and the gravity-based interaction measure (Section 3.5.2.3)—provide alternative definitions of the severity penalty. The penalty formulation, score aggregation, and margin weighting (Sections 3.5.2.4–3.5.2.5) then combine the frequency and severity contributions identically regardless of which kernel is used.

3.5.2.1 ELEMENT-LEVEL DEFECT ACCUMULATION

For a given laminate scenario $\mathbf{L}$ and defect type $m \in \mathcal{M}$, the ply-level element contribution $\delta_{m,n}(s_i)$ at element n from ply scenario s_i is retrieved from the pre-computed defect data. For area defects, $\delta_{m,n}(s_i) = \alpha_{j,n}$ is the coverage fraction from Equation (3.4), representing the proportion of element n covered by threshold-exceeding defect polygon j;

this value lies in $[0, 1]$ for each ply. For point defects, $\delta_{m,n}(s_i) = \mathbf{I}_m(d_{m,n})$ is the threshold indicator from Equation (3.7), equal to 1 if the defect value exceeds the threshold and 0 otherwise.

The accumulated defect contribution at element n across all plies is:

$$\tilde{\delta}_{m,n} = \sum_{i=1}^{N_p} \delta_{m,n}(s_i) \tag{3.45}$$

For area defects, $\tilde{\delta}_{m,n}$ represents the total defect coverage burden at element n through the laminate thickness. Values exceeding 1.0 indicate that the cumulative defect coverage at that element surpasses the element area, a signature of through-thickness stacking. For point defects, $\tilde{\delta}_{m,n}$ is the count of plies with threshold-exceeding defect values at element n. In both cases, higher values indicate greater defect concentration through the thickness. This accumulated quantity feeds the laminate-level frequency penalty defined in Section 3.5.2.4 and serves as a common input to both severity kernels.

3.5.2.2 STAGGER SHIFT SEVERITY KERNEL

The first severity kernel models through-thickness defect interaction as a proximity-weighted coincident stacking measure. The accumulation in Equation (3.45) treats all stacking equally regardless of the through-thickness separation between contributing plies. From a structural perspective, however, adjacent stacked defects are substantially more detrimental than defects separated by several well-consolidated plies [209]. Adjacent stacking creates continuous resin-rich regions or fiber-free channels that act as stress concentrations and potential delamination initiation sites, whereas intervening plies distribute load and attenuate the interaction between spatially coincident defects.

To capture this effect, the stagger shift kernel incorporates a proximity weighting that penalizes defect coincidence more heavily when the contributing plies are close together in the stacking sequence. The formulation is based on a pairwise interaction model: for each pair of plies (i, j) with $i < j$, the stacking penalty at element n is weighted by a proximity function that decays with increasing ply separation. The through-thickness separation between plies i and j is defined as:

$$\Delta_{ij} = |i - j| \tag{3.46}$$

where the indices correspond to the position of each ply in the stacking sequence. Adjacent plies have $\Delta_{ij} = 1$; plies separated by k intervening plies have $\Delta_{ij} = k + 1$. The proximity weight is then defined as an inverse-power decay:

$$w(\Delta_{ij}) = \frac{1}{\Delta_{ij}^{\beta}} \tag{3.47}$$

where $\beta > 0$ is a user-configurable decay parameter. The default value $\beta = 1$ produces inverse-linear decay, yielding the following representative weights:

Δ_{ij}	1	2	3	4	5	10
$w(\Delta_{ij})$	1.00	0.50	0.33	0.25	0.20	0.10

The choice of inverse-linear decay reflects the physical reasoning that load transfer between plies diminishes approximately inversely with separation in classical laminate theory. Setting $\beta = 0$ recovers the unweighted formulation (equal penalty for all separations), while larger values of β concentrate the penalty more sharply on adjacent plies.

The stagger shift severity at element n for defect type m is defined through the pairwise interaction term:

$$\sigma_{m,n} = \sum_{i=1}^{N_p} \sum_{j=i+1}^{N_p} w(\Delta_{ij}) \cdot \delta_{m,n}(s_i) \cdot \delta_{m,n}(s_j) \tag{3.48}$$

This term is nonzero only when two or more plies contribute defects at element n: the product $\delta_{m,n}(s_i) \cdot \delta_{m,n}(s_j)$ is zero whenever either ply is defect-free at that element. When no stacking occurs (at most one ply has a defect at element n), $\sigma_{m,n} = 0$. When both plies carry defects, the contribution is weighted by their proximity: adjacent plies ($\Delta_{ij} = 1$) contribute the full product of the two defect fractions, while distant plies contribute a reduced product scaled by $w(\Delta_{ij})$. The double sum ensures that all pairwise interactions are captured, with closer pairs receiving greater weight.

The stagger shift kernel captures direct through-thickness alignment but is insensitive to in-plane proximity. Two defects offset by a single element in-plane contribute zero to $\sigma_{m,n}$ regardless of their through-thickness separation, because the pairwise product requires both contributions to reference the same element. This limitation motivates the gravity-based kernel introduced next, which extends the interaction model to capture spatial proximity across neighboring elements.

3.5.2.3 GRAVITY-BASED SEVERITY KERNEL

The second severity kernel models through-thickness defect interaction through an analogy to Newton's law of universal gravitation [210], originally proposed by Patel for synthetic defect data on uniform rectangular grids [77]. In this formulation, each defective element is assigned a "mass" proportional to its defect density, and the interaction between two defective elements decays with the square of their combined distance in a unified 3D metric that treats in-plane surface separation and through-thickness ply separation on equivalent physical footing. The result is a severity measure that penalizes both direct stacking and near-miss proximity, providing a more comprehensive assessment of through-thickness defect arrangement than the coincident stacking kernel alone.

The formulation presented here extends Patel's approach in three respects. First, it operates on the non-uniform quadrilateral finite element mesh established in Section 3.1.1

using VCP-predicted defect data rather than synthetic rectangular defects on a uniform grid. Second, it incorporates a graph-based geodesic distance metric that respects the tool surface topology, enabling correct treatment of complex geometries such as near-radius features. Third, it integrates with the L-IDS scoring infrastructure established in Sections 3.2–3.3, including margin-of-safety weighting for the L-DIDS variant, so that the gravity-based method inherits the full scoring and margin-awareness capabilities of the framework.

Scope of Application The gravity-based severity kernel is applied exclusively to area defects (gaps and overlaps). Point defects (angle deviation and steering) retain the stagger shift coincident stacking kernel from Equation (3.48). This scope distinction reflects the physical character of the two defect categories. Gap and overlap defects are discrete, spatially localized features whose structural consequence depends on their spatial concentration and proximity; the gravitational interaction model captures this dependence naturally through the spatial decay of the kernel. Angle deviation and steering exceedances, by contrast, are distributed field quantities that produce broad contiguous regions of threshold exceedance; their structural consequence is better characterized by the coincident stacking count at each element than by a spatially-decaying interaction model.

The per-defect-type penalty structure of the L-IDS framework (Section 3.2.6) supports this hybrid treatment without modification. Each defect type m has its own severity penalty $P^{\text{sev}}_{m,\text{lam}}$ computed independently. For m = gap and m = overlap, the gravity-based kernel produces this penalty. For m = angle deviation and m = steering, the stagger shift kernel produces it. The downstream AHP aggregation combines all defect types identically regardless of which kernel produced the severity.

Distance Metric The gravitational interaction kernel requires a distance between any two elements that combines their in-plane surface separation with their through-thickness separation in the stacking sequence. Both components are defined in physical units (inches) to ensure dimensional consistency.

The in-plane distance B_{nk} between elements n and k is defined as the approximate geodesic distance along the tool surface, computed as the shortest-path distance on the finite element mesh graph. The mesh graph is constructed by assigning each element a node and connecting two elements with an edge if and only if they share exactly two mesh nodes (i.e., they share a mesh edge). The weight assigned to each graph edge is the Euclidean distance between the 3D centroids of the connected elements:

$$w_{\text{edge}}(n, k) = \|\mathbf{c}_n - \mathbf{c}_k\| \tag{3.49}$$

where $\mathbf{c}_n = (x_n, y_n, z_n)$ and $\mathbf{c}_k$ are the physical centroid coordinates stored in the mesh data structure (Section 3.1.1). The shortest-path distance from element n to element k in this weighted graph approximates the geodesic distance along the tool surface.

This graph-based distance has three properties relevant to the gravity-based formulation. It produces distances in physical units (inches), consistent with tow width and ply thickness. It respects the surface topology: elements that are close in 3D Euclidean distance but far apart along the surface—such as elements on opposite sides of a tight-radius feature—receive a large graph distance reflecting their true along-surface separation. Finally, it is independent of the NURBS parametric coordinates of the tool surface, ensuring that the optimization result does not depend on the surface parameterization.

The through-thickness distance between plies i and j is defined using a normalized metric that places the through-thickness and in-plane dimensions on comparable physical scales:

$$\bar{z}_{ij} = |i - j| \cdot w_{\text{tow}} \tag{3.50}$$

where $|i - j|$ is the ply index separation and w_{tow} is the tow width. The normalization by tow width rather than by physical ply thickness addresses a fundamental scale mismatch in composite laminates. Ply thickness (typically 0.0086 in. for carbon fiber prepreg) is roughly 30 to 60 times smaller than tow width (typically 0.25–0.50 in.). Without normalization, the squared distance denominator in the gravitational kernel would be dominated by the in-plane term for any non-stacked element pair, rendering the proximity contribution negligible and collapsing the gravity-based kernel's behavior to approximately that of the coincident stacking kernel. The normalization ensures that one ply of through-thickness separation is treated as comparable in importance to one tow width of in-plane separation, which is appropriate because both represent one unit of the fundamental length scale governing AFP defect geometry.

The total squared distance between element n in ply i and element k in ply j is then:

$$d_{nk,ij}^2 = B_{nk}^2 + \bar{z}_{ij}^2 \tag{3.51}$$

which treats the in-plane surface distance and the normalized through-thickness separation as orthogonal components of a combined defect separation metric.

Interaction Kernel The gravitational interaction kernel defines the strength of the interaction between two defective elements as a function of their combined distance. Following the point-mass approximation of Newton's inverse-square law, the kernel is:

$$\mathcal{G}(n, k, \bar{z}_{ij}) = \frac{G}{B_{nk}^2 + \bar{z}_{ij}^2} \tag{3.52}$$

where G is the gravitational normalization constant defined below. The point-mass approximation treats each element as a point source located at its centroid. This approximation introduces error proportional to (element size/separation distance)^2, which is negligible for all element pairs except those at identical spatial locations across plies. For such directly stacked elements ($B_{nk} = 0$), the kernel reduces to $G/\bar{z}_{ij}^2$, which is well-defined and finite for any pair of distinct plies.

The constant G is chosen so that the maximum pairwise interaction between two fully defective, directly stacked, adjacent-ply elements equals unity:

$$G = w_{\text{tow}}^2 \tag{3.53}$$

This normalization is verified by evaluating the kernel at $B_{nk} = 0$ and $\bar{z}_{ij} = w_{\text{tow}}$ (adjacent plies):

$$\mathcal{G} = \frac{w_{\text{tow}}^2}{0 + w_{\text{tow}}^2} = 1 \tag{3.54}$$

All other element pairs produce kernel values strictly less than 1, and elements further apart in-plane or separated by more plies produce values that decay toward zero. This normalization ensures that the gravity-based severity is scaled consistently with the stagger shift severity so that both kernels produce scores on the same L-IDS interval.

Interaction Cutoff The gravitational kernel decays rapidly with distance. For adjacent plies, the kernel value at an in-plane separation of r tow widths is $1/(r^2 + 1)$, which yields 0.500 at $r = 1$, 0.100 at $r = 3$, 0.015 at $r = 8$, and 0.004 at $r = 15$. Beyond one course width ($r = N_{\text{tows}}$), the interaction is below 2% of its maximum and contributes negligibly to the severity. A cutoff radius is therefore applied to limit the interaction sum to physically meaningful element pairs.

The natural physical bound for the cutoff is one course width, beyond which defect interactions are structurally negligible. This bound is sufficient for typical AFP applications where the mesh density is fine relative to the course width. However, the gravity-based kernel's distinguishing capability—its sensitivity to in-plane proximity—requires that the cutoff radius encompass enough neighboring elements to express a meaningful spatial decay structure at the local mesh resolution. On meshes that are coarse relative to the course width, the course-width bound alone may capture only a small number of immediate edge-adjacent neighbors, in which case the gravity-based kernel reduces in practice toward the coincident-stacking behavior of the stagger shift kernel because too few near-miss neighbors exist within the cutoff to influence the optimization. To ensure consistent kernel behavior across the full range of plausible mesh densities, a mesh-relative floor is added to the cutoff definition.

The cutoff radius is therefore defined as the larger of two characteristic length scales:

$$r_{\text{cutoff}} = \max\left(N_{\text{tows}} \cdot w_{\text{tow}},\ k \cdot \sqrt{\bar{A}_e}\right) \tag{3.55}$$

where N_{tows} is the number of tows per course, $\bar{A}_e$ is the mean element area across the mesh, $\sqrt{\bar{A}_e}$ is therefore the characteristic element side length, and k is a fixed multiplier set to $k = 4$. The first argument is the physical course-width bound; the second argument is a mesh-relative floor that guarantees a minimum reach of four element side lengths, which is sufficient to encompass roughly two rings of in-plane neighbors around each reference element. The choice $k = 4$ corresponds to the minimum reach at which the gravity-based kernel can resolve a meaningful proximity decay structure: it captures the self-interaction term, the immediate edge-adjacent neighbors, and one additional ring of neighbors beyond. Smaller values of k risk collapsing the kernel toward coincident-stacking-only behavior on coarse meshes; larger values increase the neighbor count without proportional benefit because the kernel decay drives distant contributions toward zero regardless.

In typical applications, the physical course-width bound dominates and the mesh-relative floor has no effect on the cutoff. The mesh-relative floor functions as a robustness measure that maintains the gravity-based kernel's spatial sensitivity in cases where the mesh density is unusually coarse relative to the course width. This ensures that the kernel's distinguishing behavior is preserved across mesh configurations without requiring users to manually retune the cutoff for each new geometry.

For each element n, the set of neighboring elements within the cutoff radius is defined as:

$$\mathcal{N}(n) = \{\, k \;:\; B_{nk} \leq r_{\text{cutoff}} \,\} \tag{3.56}$$

Because the mesh geometry does not change between laminate evaluations, $\mathcal{N}(n)$ and the associated distances B_{nk} are precomputed once before the optimization loop begins. The precomputation uses bounded Dijkstra shortest-path searches on the mesh graph: starting from each element n, the search expands outward along mesh edges, accumulating path distances, and terminates when the accumulated distance exceeds r_{cutoff}. The result is a sparse neighbor list with associated geodesic distances for each element, stored as a precomputed data structure indexed by element ID.

Severity Formulation The gravity-based severity at element n for area defect type m is defined as:

$$F_{m,n}^{\text{grav}} = \frac{1}{2} \sum_{i=1}^{N_p} \sum_{j=i+1}^{N_p} \sum_{k \in \mathcal{N}(n)} \mathcal{G}(n, k, \bar{z}_{ij}) \cdot \Big[\delta_{m,n}(s_i) \cdot \delta_{m,k}(s_j) + \delta_{m,n}(s_j) \cdot \delta_{m,k}(s_i) \Big] \tag{3.57}$$

The three summation indices correspond to three distinct contributions. The outer double sum over ply pairs (i, j) with $i < j$ enumerates all unique through-thickness ply pairings, consistent with the pairwise structure of the stagger shift severity in Equation (3.48). The inner sum over k enumerates all elements within the precomputed neighborhood of element n, including element n itself ($k = n$, the self-interaction case corresponding to direct stacking).

For each ply pair and each neighbor element, two interaction terms are computed.

The first term, $\delta_{m,n}(s_i) \cdot \delta_{m,k}(s_j)$, captures the interaction of element n in ply i with element k in ply j. The second term, $\delta_{m,n}(s_j) \cdot \delta_{m,k}(s_i)$, captures the interaction of element n in ply j with element k in ply i. Both terms are required because element n may have different defect contributions in different plies. The leading factor of $1/2$ corrects for this symmetric double-counting, so that the maximum severity contribution from a single fully defective stacked pair between two adjacent plies equals unity.

The formulation exhibits three properties that confirm its physical consistency. When element n is defect-free in all plies, every term in the sum contains a zero factor, so $F^{\text{grav}}_{m,n} = 0$; defect-free elements contribute no severity. When at most one ply contributes a defect at or near element n, the pairwise product requires two simultaneously nonzero contributions and cannot be satisfied, so $F^{\text{grav}}_{m,n} = 0$ as well; isolated defects with nothing to interact with produce no severity. Finally, when the neighbor sum is restricted to the self-interaction term $k = n$ only, substitution of $\mathcal{G}(n, n, \bar{z}_{ij}) = G/\bar{z}^2_{ij} = 1/|i - j|^2$ reduces the gravity-based severity to the stagger shift form with $\beta = 2$:

$$F^{\text{grav}}_{m,n}\Big|_{k=n} = \sum_{i=1}^{N_p} \sum_{j=i+1}^{N_p} \frac{1}{|i - j|^2} \cdot \delta_{m,n}(s_i) \cdot \delta_{m,n}(s_j) \tag{3.58}$$

demonstrating that the gravity-based kernel is a strict generalization of the stagger shift kernel with a specific inverse-square through-thickness decay: the additional contributions from non-self neighbors ($k \neq n$) are what distinguish the two kernels and give the gravity-based method its sensitivity to in-plane proximity.

3.5.2.4 LAMINATE-LEVEL PENALTY FORMULATION

The laminate-level penalties follow the frequency-severity structure established at the ply level in Section 3.2.5, with two modifications: the frequency penalty uses the accumulated defect contributions across all plies, and the severity penalty is computed using one of the two severity kernels defined above. The frequency penalty and the penalty aggregation pipeline are identical for both kernels; only the severity term differs.

The laminate-level manufacturing frequency penalty for defect type m measures the total defect burden across all plies, normalized by the laminate boundary area:

$$P_{m,\mathrm{lam}}^{\mathrm{freq}} = \frac{1}{A_{\mathrm{lam}}} \sum_{n=1}^{N_e} \tilde{\delta}_{m,n} \cdot A_n \tag{3.59}$$

where A_{lam} is the total laminate boundary area (the union of all ply boundaries) and A_n is the area of element n. The frequency penalty captures the spatial extent of defects across the laminate without regard to their through-thickness distribution. It can exceed 1.0 when multiple plies contribute defects at the same elements, which is physically meaningful: a value of 1.5 indicates that defects cover, in aggregate, 150% of the laminate area through the thickness.

The margin-weighted variant replaces the manufacturing contributions with margin-weighted contributions following Section 3.3.4:

$$P_{m,\mathrm{lam,margin}}^{\mathrm{freq}} = \frac{1}{A_{\mathrm{lam,margin}}} \sum_{n=1}^{N_e} \frac{\tilde{\delta}_{m,n} \cdot A_n}{\mathrm{FoS}_n} \tag{3.60}$$

where $A_{\mathrm{lam,margin}} = \sum_{n=1}^{N_e} A_n/\mathrm{FoS}_n$ is the margin-weighted laminate area.

The normalization by A_{lam} rather than by the sum of individual ply areas $\sum_i A_{\mathrm{ply},i}$ is a deliberate choice. The laminate boundary area represents the physical footprint over which the laminate exists, and normalization by this footprint produces a score that reflects the local concentration of defects at each element through the thickness. This is the appropriate normalization for a metric targeting stacking severity, as it weights elements by their contribution to the structural cross-section rather than by the total material placed.

The laminate-level manufacturing severity penalty for defect type m is defined using the selected severity kernel. Let $\phi_{m,n}$ denote the element-level severity, which takes one of two forms depending on the kernel selection and defect type:

$$\phi_{m,n} = \begin{cases} \sigma_{m,n} & \text{if stagger shift kernel is selected, or } m \in \{\text{angle dev., steering}\} \\ F^{\text{grav}}_{m,n} & \text{if gravity-based kernel is selected and } m \in \{\text{gap, overlap}\} \end{cases} \tag{3.61}$$

The manufacturing severity penalty is then:

$$P^{\text{sev}}_{m,\text{lam}} = \frac{1}{A_{\text{lam}}} \sum_{n=1}^{N_e} \phi_{m,n} \cdot A_n \tag{3.62}$$

where $\phi_{m,n}$ is supplied by Equation (3.48) for the stagger shift kernel or by Equation (3.57) for the gravity-based kernel. In both cases, the severity penalty is zero when no through-thickness defect interaction occurs and increases with the spatial extent and concentration of interacting defects.

The distinction between the ply-level and laminate-level severity formulations reflects their different physical targets. At the ply level, severity measures how badly individual defects exceed their thresholds (the quadratic function $s_m(d) = (d/\tau_m)^2$ from Equation (3.9)). At the laminate level, severity measures how densely defects are concentrated through the thickness, with the specific form of the density measure determined by the selected kernel. The two formulations are architecturally parallel—both multiply a spatial contribution by an amplification factor—but address distinct aspects of manufacturing quality.

The margin-weighted severity penalty follows the same pattern:

$$P^{\text{sev}}_{m,\text{lam,margin}} = \frac{1}{A_{\text{lam,margin}}} \sum_{n=1}^{N_e} \frac{\phi_{m,n} \cdot A_n}{\text{FoS}_n} \tag{3.63}$$

The margin weighting is applied at the penalty summation level rather than within the severity kernel itself. This placement maintains exact architectural parallelism with the ply-level DIDS formulation (Section 3.3.4) and produces the intended structural-awareness behavior: severity contributions from elements in structurally critical regions (low FoS_n) are amplified in the penalty sum, so the optimizer is driven to disperse defects away from those regions. The same weighting applies to both severity kernels, ensuring that the gravity-based and stagger shift variants produce comparable L-DIDS scores that inherit the full margin-awareness capability of the framework.

3.5.2.5 SCORE AGGREGATION

The laminate-level penalties are aggregated following the same procedure used at the ply level in Section 3.2.6. For each defect type m, the weighted penalty combines the frequency and severity components:

$$P_{m,\mathrm{lam}} = \omega_m^{\mathrm{freq}} \cdot P_{m,\mathrm{lam}}^{\mathrm{freq}} + \omega_m^{\mathrm{sev}} \cdot P_{m,\mathrm{lam}}^{\mathrm{sev}} \tag{3.64}$$

The total manufacturing penalty is the sum over all defect types:

$$\mathrm{Penalty}_{\mathrm{lam,mfg}} = \sum_{m \in \mathcal{M}} P_{m,\mathrm{lam}} \tag{3.65}$$

The L-MIDS is then defined as:

$$\text{L-MIDS} = 100 - \mathrm{Penalty}_{\mathrm{lam,mfg}} \tag{3.66}$$

The L-DIDS is computed analogously using the margin-weighted penalties from Equations (3.60) and (3.63):

$$\text{L-DIDS} = 100 - \mathrm{Penalty}_{\mathrm{lam,margin}} \tag{3.67}$$

The L-IDS combines both perspectives into a single objective:

$$\text{L-IDS}(\mathbf{L}) = w_{\text{design}} \cdot \text{L-DIDS} + w_{\text{mfg}} \cdot \text{L-MIDS} \tag{3.68}$$

where w_{design} and w_{mfg} are the user-defined weighting factors satisfying $w_{\text{design}} + w_{\text{mfg}} = 1$, consistent with the IDS formulation in Equation (3.29).

The laminate-level score is bounded above by 100, achieved when no defects in any ply exceed their respective thresholds. The theoretical lower bound differs from zero because the accumulated frequency penalties can exceed 1.0 when multiple plies contribute defects at the same elements, and both severity kernels can produce per-element contributions exceeding unity under heavy stacking conditions. Laminate scores below approximately 50 indicate substantial through-thickness defect concentration, and scores near or below zero indicate conditions requiring fundamental revision of the process plan. The precise interpretation depends on the laminate configuration and the selected severity kernel, as documented in Section 4.5. The extended dynamic range relative to the ply-level score is a deliberate feature of the laminate formulation, providing greater discrimination between moderate and severe stacking conditions that could be indistinguishable under a score clamped to $[0, 100]$.

The two severity kernels produce scores on the same scale, bounded above by 100, but are not numerically equivalent: the gravity-based kernel generally produces more penalty per unit of stacking because it incorporates spatially-decaying interactions that the stagger shift kernel does not capture. A gravity-based L-IDS of 85 therefore reflects a different assessment than a stagger shift L-IDS of 85 for the same laminate, even though both are on the same bounded interval. The two kernels are intended to be used alternatively rather than simultaneously, with the user selecting the kernel whose modeling assumptions best match the application's structural concerns. Direct numerical comparison between the two kernel outputs is presented in Section 4.5 to illustrate their differing sensitivities.

3.5.3 OPTIMIZATION ALGORITHM

The combinatorial complexity established in Equation (3.44) precludes exhaustive search. The LLO algorithm addresses this through cyclic coordinate optimization—a strategy that decomposes the multi-variable combinatorial problem into a sequence of single-variable subproblems by optimizing one ply at a time while fixing all others [211]. Multiple passes through the ply sequence are executed until a fixed-point convergence criterion is satisfied. The same algorithm is used for both the stagger shift and gravity-based variants; only the severity kernel inside the L-IDS evaluation differs.

3.5.3.1 ALGORITHM STRUCTURE

Let $\mathbf{L}^{(t)}$ denote the laminate configuration at the end of pass t, and let $\mathbf{L}^{(0)}$ denote the baseline configuration. Within pass t, the algorithm processes plies in stacking sequence order ($i = 1, 2, \ldots, N_p$). For each ply i, the algorithm evaluates every candidate scenario $s' \in \mathcal{P}_i$ by constructing a trial laminate that substitutes s' for the current assignment of ply i while retaining all other assignments:

$$\mathbf{L}_{\text{trial}} = \left(s_1^{(t)},\ \ldots,\ s_{i-1}^{(t)},\ s',\ s_{i+1}^{(t-1)},\ \ldots,\ s_{N_p}^{(t-1)}\right) \tag{3.69}$$

The notation reflects the Gauss-Seidel update structure: plies 1 through $i-1$ use their already-updated assignments from the current pass, while plies $i+1$ through N_p retain their assignments from the previous pass. Each trial laminate is scored using the full laminate-level objective L-IDS, evaluating defect stacking across all N_p plies. The candidate that maximizes the laminate score is selected:

$$s_i^{(t)} = \underset{s' \in \mathcal{P}_i}{\arg\max}\ \text{L-IDS}(\mathbf{L}_{\text{trial}}) \tag{3.70}$$

When multiple candidates achieve the same maximum laminate score, the tie is broken by selecting the candidate with the highest ply-level IDS, favoring the scenario that offers the best individual-ply quality among equally performing laminate configurations.

The complete algorithm is presented in Algorithm 3.

Algorithm 3 Laminate-Level Optimization

Require: Baseline laminate $\mathbf{L}^{(0)}$, candidate pools $\{\mathcal{P}_1, \ldots, \mathcal{P}_{N_p}\}$, severity kernel selection (stagger shift or gravity-based), maximum passes $T_{\max}$
Ensure: Optimized laminate $\mathbf{L}^*$

if gravity-based kernel is selected **then**
    Precompute neighbor sets $\mathcal{N}(n)$ and geodesic distances B_{nk} via bounded Dijkstra
end if

for $t = 1, 2, \ldots, T_{\max}$ **do**
    $\mathbf{L}^{(t)} \leftarrow \mathbf{L}^{(t-1)}$ ▷ Initialize pass from previous result
    for $i = 1, 2, \ldots, N_p$ **do** ▷ Stacking sequence order
        $s_{\text{best}} \leftarrow s_i^{(t)}, \quad f_{\text{best}} \leftarrow \text{L-IDS}(\mathbf{L}^{(t)})$
        for each $s' \in \mathcal{P}_i$ **do**
            Construct $\mathbf{L}_{\text{trial}}$ per Equation (3.69)
            $f' \leftarrow \text{L-IDS}(\mathbf{L}_{\text{trial}})$
            if $f' > f_{\text{best}}$ **or** $\big(f' = f_{\text{best}}$ **and** $\text{IDS}(s') > \text{IDS}(s_{\text{best}})\big)$ **then**
                $s_{\text{best}} \leftarrow s', \quad f_{\text{best}} \leftarrow f'$
            end if
        end for
        $s_i^{(t)} \leftarrow s_{\text{best}}$
    end for

    if $\mathbf{L}^{(t)} = \mathbf{L}^{(t-1)}$ **then** ▷ Ply scenario stability
        break ▷ Converged
    end if
end for
return $\mathbf{L}^* \leftarrow \mathbf{L}^{(t)}$

3.5.3.2 CONVERGENCE PROPERTIES

The algorithm possesses three properties that guarantee finite convergence. First, monotonicity: each ply update in Equation (3.70) selects the candidate that maximizes L-IDS over $\mathcal{P}_i$. Because the current assignment $s_i^{(t)}$ is always a member of the candidate pool (the baseline scenario is always included, and the current assignment is either the baseline or a previously selected candidate), the score after updating ply i is at least as good as the score before the update. This holds for every ply within a pass, establishing that $\text{L-IDS}(\mathbf{L}^{(t)}) \geq \text{L-IDS}(\mathbf{L}^{(t-1)})$ for all t. Second, boundedness: the score is bounded above by 100, corresponding to the defect-free laminate. Third, finite convergence: the candidate pools are finite, so the set of all possible laminate configurations Ω is finite. The score sequence $\{\text{L-IDS}(\mathbf{L}^{(t)})\}_{t=0}^{\infty}$ is non-decreasing and bounded above. A non-decreasing bounded sequence on a finite set must become constant after finitely many steps. Because each pass contains N_p ply updates and each update selects from at most $\max_i |\mathcal{P}_i|$ candidates, the algorithm terminates within a finite number of passes, as confirmed in the application results presented in Section 4.5.

These convergence properties are standard for cyclic coordinate optimization on finite discrete domains [211]. The guarantee applies regardless of ply processing order, initialization, or the specific form of the scoring function, requiring only monotonicity and finiteness. Critically, none of these properties depend on the internal structure of the severity kernel, so the convergence guarantee applies equally to the stagger shift and gravity-based variants.

At convergence, the solution $\mathbf{L}^*$ is a *coordinate-wise optimum*: no single-ply substitution can improve the laminate score. Formally, for every ply i and every candidate $s' \in \mathcal{P}_i$:

$$\text{L-IDS}(s_1^*, \ldots, s_{i-1}^*, s', s_{i+1}^*, \ldots, s_{N_p}^*) \leq \text{L-IDS}(\mathbf{L}^*) \tag{3.71}$$

This condition is analogous to a Nash equilibrium in the ply assignments: each ply's selection is optimal given all other plies' selections. A coordinate-wise optimum is not guaranteed to be a global optimum; escaping it would require simultaneous multi-ply substitutions. However, for the LLO problem, the dominant stacking interactions involve pairs of plies with identical fiber orientations, and the coordinate-wise method effectively addresses these pairwise interactions through sequential refinement with multi-pass propagation.

3.5.4 ALGORITHM SELECTION RATIONALE

The choice of cyclic coordinate optimization over alternative approaches—particularly population-based metaheuristics such as genetic algorithms—is motivated by the specific structure of the LLO problem.

The LLO design variables exhibit a decomposable structure: each ply's defect distribution is determined independently by VCP, and the interaction between plies occurs only through spatial overlap of their defect geometries on the shared finite element mesh. This decomposability is directly exploited by the coordinate-wise method, which reduces the combinatorial problem to a sequence of tractable single-ply enumerations. Metaheuristic methods do not exploit this structure; they treat the configuration space as monolithic and rely on stochastic exploration, which is less efficient for problems with natural coordinate separability.

The computational bottleneck of the LLO workflow is VCP simulation, which requires minutes per ply scenario. The coordinate-wise method confines all VCP calls to a one-time candidate pool generation phase, after which the iterative optimization operates entirely on pre-computed defect data. A genetic algorithm exploring novel ply scenarios beyond the pre-computed pool would require thousands of VCP evaluations—an order-of-magnitude increase in computational cost. Even when restricted to the pre-computed pool, a genetic algorithm with a population size of 50 running for 100 generations would require 5,000 laminate score evaluations per generation for crossover and mutation assessment, compared to approximately 77 evaluations per pass in the coordinate method for an 11-ply laminate with 7 candidates per ply.

The coordinate-wise method provides provable convergence to a coordinate-wise optimum in finite iterations (Section 3.5.3.2). Metaheuristic methods offer no analogous convergence guarantee; their performance depends on parameter tuning (mutation rate, crossover operators, selection pressure) that is problem-specific and requires empirical calibration. For an optimization framework embedded within an engineering tool, deterministic convergence with predictable runtime is preferable to stochastic search with tunable but unpredictable behavior.

The gap between the coordinate-wise optimum and the global optimum could, in principle, be explored through multi-ply swap neighborhoods or hybridization with a metaheuristic refinement stage. The modular architecture of CAPP accommodates such extensions: the coordinate-wise method could serve as an efficient initialization for a subsequent metaheuristic phase. This hybrid approach represents a natural direction for future development should computational resources or problem scale warrant it, and is discussed further in Section 5.3.

3.5.5 SCOPE OF LAMINATE-LEVEL DEFECT INTERACTIONS

The two severity kernels presented in Section 3.5.2 capture different aspects of laminate-level defect interaction, and each has its own scope limitations with respect to defect-type combinations and interaction modes.

Both kernels compute defect accumulation and stacking severity independently for each defect type. Same-type stacking—gap-on-gap and overlap-on-overlap coincidence through the thickness—is therefore captured explicitly by both kernels. These are the stacking configurations of greatest structural consequence in AFP manufacturing, because plies with identical fiber orientations produce geometrically similar defect patterns that stack directly when starting points are not staggered. The stagger shift kernel captures this same-type stacking through coincident pairwise products at each element. The gravity-based kernel captures it through spatially-decaying pairwise products across local neighborhoods, with the self-interaction term $k = n$ reproducing the coincident stacking behavior and the off-diagonal terms adding proximity sensitivity.

Cross-type defect interactions, in which a gap in one ply coincides spatially with an overlap in another ply (or with an angle deviation exceedance), are not captured by either kernel in the current formulation. This is a deliberate modeling decision motivated by two considerations. First, the AFP research community has not established quantitative relationships for how different cross-type interactions affect structural performance. The relative severity of gap-gap versus overlap-overlap versus gap-overlap stacking is not characterized by existing experimental or analytical studies, and any weighting scheme for cross-type interactions would therefore be speculative. Second, the relative importance of cross-type interactions is application-dependent: a part with a mating surface may be more sensitive to overlap-overlap stacking, which alters surface profile, while a structurally loaded panel may be more sensitive to gap-gap stacking, which creates resin-rich channels.

The existing defect weight system (Section 3.2.1) already provides a mechanism for users to prioritize same-type stacking by adjusting the relative weights of gap versus overlap penalties. This user-level customization accommodates application-specific concerns within the current framework.

The architecture of the laminate-level scoring is designed to be extensible to cross-type interactions under either severity kernel. An interaction penalty term could be added as an additional defect type within the existing $\mathcal{M}$ framework, with its own threshold, weights, and penalty computation following the modular structure used for gaps, overlaps, angle deviation, and steering. Specifically, the pairwise formulations in Equations (3.48) and (3.57) could be extended to compute cross-type products $\delta_{m_1,n}(s_i)\cdot\delta_{m_2,k}(s_j)$ for $m_1 \neq m_2$, with interaction-specific weights. As experimental or simulation-based studies establish quantitative relationships between cross-type stacking and structural knockdown, these relationships can be incorporated without restructuring the optimization algorithm.

A second scope distinction concerns the treatment of point defects under the gravity-based kernel. As noted in Section 3.5.2.3, the gravity-based kernel is applied only to area defects (gaps and overlaps); point defects (angle deviation and steering) retain the stagger shift kernel regardless of the user's overall kernel selection. This hybrid treatment reflects the different physical character of the two defect categories: area defects are discrete localized features for which the gravitational proximity model is physically meaningful, while point defects are distributed field quantities for which coincident stacking is the more appropriate measure. Users who select the gravity-based kernel therefore obtain a mixed scoring configuration in which area defect severity is computed gravitationally and point defect severity is computed coincidentally, with the downstream AHP aggregation combining both into a single L-IDS value.

3.5.6 COMPUTATIONAL ANALYSIS

The computational cost of LLO consists of three distinct phases: candidate pool generation (VCP-bound), neighbor precomputation (required only for the gravity-based kernel), and iterative optimization (scoring-bound). The first and third phases are common to both severity kernels; the second is an additional one-time cost incurred only when the gravity-based kernel is selected.

3.5.6.1 CANDIDATE POOL GENERATION

In Mode A, the number of VCP evaluations required to populate the candidate pools is:

$$N_{\text{VCP}} = \sum_{i=1}^{N_p} \left(|\mathcal{P}_i| - 1 \right) \tag{3.72}$$

where the -1 accounts for the baseline scenario, which is already processed. For an 11-ply laminate with 7 candidates per ply (6 new shifts per ply), this yields $N_{\text{VCP}} = 66$ evaluations. Each evaluation requires a VCP simulation whose duration depends on part size, mesh density, and course width. Using the subprocess integration described in Section 3.1.6, individual ply regenerations require approximately 15–30 seconds for typical geometries, placing the total pool generation time at 15–35 minutes.

In Mode B, the candidate pool is inherited from PLO and no additional VCP evaluations are required.

Because both severity kernels operate on the same pre-computed defect data, the candidate pool generation cost is identical whether the user selects the stagger shift or gravity-based variant.

3.5.6.2 NEIGHBOR PRECOMPUTATION (GRAVITY-BASED KERNEL)

When the gravity-based kernel is selected, the neighbor sets $\mathcal{N}(n)$ and associated geodesic distances B_{nk} must be computed once before the optimization loop. The precomputation consists of two steps: constructing the mesh graph by identifying edge-adjacent element pairs from the nodal connectivity, and executing bounded Dijkstra shortest-path searches from each element to populate the neighbor list.

The mesh graph construction requires a single pass through the element list to build a node-to-element index, followed by a pass through each element's four corner nodes to identify neighbors sharing exactly two mesh nodes. For a mesh with N_e elements, this operation is $O(N_e)$. The bounded Dijkstra from each element visits only elements within the cutoff radius; its cost is $O(N_{\text{neighbors}} \log N_{\text{neighbors}})$ per element, where $N_{\text{neighbors}}$ is the average number of elements within r_{cutoff}. The total precomputation cost is therefore $O(N_e \cdot N_{\text{neighbors}} \log N_{\text{neighbors}})$.

For the Quartic benchmark mesh ($N_e = 3{,}645$, mean element area $\bar{A}_e \approx 3.03$ sq in., mean element side length $\sqrt{\bar{A}_e} \approx 1.74$ in.) with the cutoff radius defined by Equation (3.55), the mesh-based reach $k \cdot \sqrt{\bar{A}_e} \approx 6.96$ in. exceeds the course-width bound (4 in. for the eight 0.5-in. tows used in the Quartic configuration), so the cutoff is governed by the mesh resolution. The precomputation produces an average of approximately 32 neighbors per element (median 31, maximum 43) and completes in under one tenth of a second on standard hardware. This cost is negligible relative to the VCP candidate generation and is incurred only once per optimization run.

3.5.6.3 ITERATIVE OPTIMIZATION

Each pass of the coordinate-wise algorithm evaluates:

$$N_{\text{score/pass}} = \sum_{i=1}^{N_p} |\mathcal{P}_i| \tag{3.73}$$

laminate score computations. For an 11-ply laminate with 7 candidates per ply, this is 77 evaluations per pass. Each laminate score evaluation involves element-level accumulation across all plies and severity computation. With convergence typically occurring within $T = 2$–4 passes, the total scoring cost is:

$$N_{\text{score,total}} = T \cdot \sum_{i=1}^{N_p} |\mathcal{P}_i| \tag{3.74}$$

yielding approximately 150–310 score evaluations for this configuration.

The cost per score evaluation depends on the severity kernel. For the stagger shift kernel, the severity computation is $O(N_e \cdot N_p^2)$ because the pairwise ply sum at each element evaluates $N_p(N_p - 1)/2$ terms. For the gravity-based kernel, the severity computation is $O(N_e \cdot N_{\text{neighbors}} \cdot N_p^2)$ because the pairwise ply sum is extended to include the neighbor sum within each element's precomputed neighborhood. The additional factor of $N_{\text{neighbors}}$ reflects the cost of the spatial summation over the neighbor set.

For the Quartic configuration, the gravity-based kernel's per-evaluation cost is higher than the stagger shift kernel's by a factor of $N_{\text{neighbors}} \approx 32$, but each kernel operation is a single arithmetic expression on precomputed quantities. With $N_e = 3{,}645$, $N_p = 11$, and $N_{\text{neighbors}} \approx 32$, the gravity-based severity computation for a single defect type involves on the order of $N_e \cdot N_{\text{neighbors}} \cdot N_p(N_p - 1)/2 \approx 6.4 \times 10^6$ kernel evaluations per laminate, completing in well under a second per laminate evaluation. The total optimization time is therefore dominated by the VCP candidate generation phase under either kernel, with the iterative optimization adding minimal overhead. This separation of concerns—expensive simulation performed once, followed by fast iterative optimization—is a key architectural advantage of the approach. Empirical confirmation of these computational characteristics, the convergence behavior, and the comparative performance of the two severity kernels on the Quartic laminate geometry is presented in Section 4.5.

3.6 INCORPORATION OF INSPECTION DATA

Once the AFP machine executes the NC programs to fabricate a laminate, both planned and unplanned Manufacturing Induced Properties (MIPs) or defects inevitably arise. In current practice, AFP operators typically perform manual inspection after each ply or during layup to identify and rework defects before proceeding to the next ply. However, because the AFP machine must be locked out during this process, manual inspection introduces significant increases in manufacturing time and cost [26, 27, 135]. In response to these limitations, the composites industry has placed considerable emphasis on automating AFP inspection to increase throughput and improve defect traceability [41, 142, 212].

Profilometers are commonly used in automated AFP inspection to detect ply-level defects such as gaps, overlaps, angle deviations, and tow misalignments. Electroimpact's RIPITx is one such system, integrated directly onto the AFP machine head [59]. The system captures high-resolution, in-situ surface data using laser profilometry as each ply is laid down, as illustrated in Figure 3.7. Raw measurement data is collected synchronously during the layup process, generating a continuous stream of positional and geometric information. Post-production, RIPITx transforms this data into point-wise defect records along each course, describing the as-manufactured conditions of each ply.

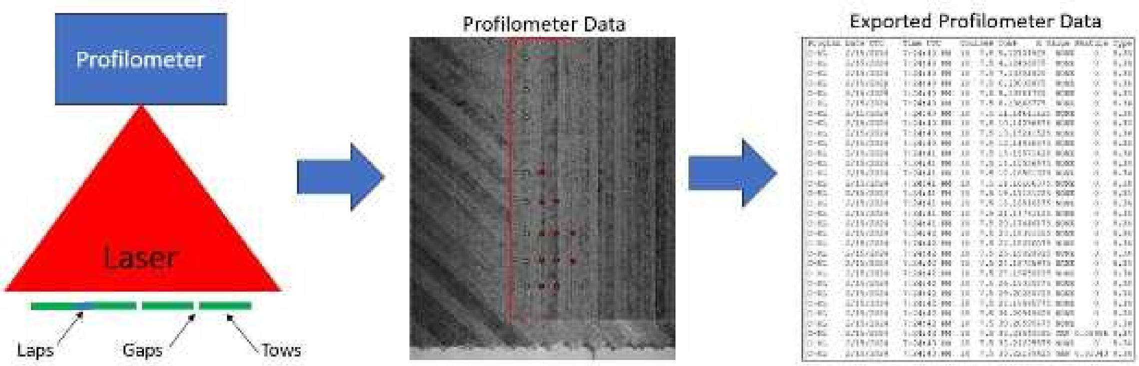

Figure 3.7: Overview of the RIPITx inspection system, showing the data acquisition workflow for defect detection [59].

This section presents the integration of RIPITx inspection data within the iSPP framework. By incorporating measured defect data directly into CAPP, engineers can visualize and compare as-manufactured defects with the as-designed predictions generated by VCP. This closed-loop capability transforms inspection data from an isolated quality-control artifact into an active component of the AFP process planning cycle, enabling identification of recurring defect patterns and informed refinement of process parameters. Demonstration of this capability on the Quartic benchmark surface is presented in Section 4.6.

3.6.1 RIPITX DATA STRUCTURE

For each AFP course, the RIPITx system generates Tab-Separated Values (TSV) files that document detected defects and corresponding course endpoints. Each ply's inspection dataset includes a defect file that records the type of defect encountered (gap or overlap), its location and orientation in three-dimensional space, and its measured width. A companion file documents the start and stop locations for each course. Together, these files provide a complete record of defect occurrences observed during production of a given ply.

The defects reported by RIPITx represent all gaps and overlaps detected during layup, including both intra-course defects arising from tows that erroneously wander from or overlap with their neighboring tows within a course, and inter-course defects occurring at the boundaries between adjacent courses. For path point locations where not all tows are active, the inspection range is reduced accordingly.

This detection scope establishes an important relationship between RIPITx and VCP. While VCP predicts inter-course defects—gaps and overlaps arising from geometric interactions between neighboring courses—RIPITx captures both inter-course defects observed during layup and intra-course defects arising from tow behavior within a single course, such as localized tow-to-tow gaps. The intra-course defects detected by RIPITx are not predicted by VCP, making the two data sources complementary: VCP provides a geometric prediction of inter-course defect distributions prior to manufacturing, while RIPITx provides a complete measured record of both defect classes as actually produced. Integrating both within CAPP therefore provides a more comprehensive view of the as-manufactured state than either source alone.

3.6.2 DATA IMPORT AND COORDINATE TRANSFORMATION

CAPP is configured to import the RIPITx TSV data files and display them alongside the predicted course information to facilitate direct comparison. The import process is managed through a dedicated dialog within the CAPP interface (Figure 3.8), where the user specifies the location of the inspection data files for either the entire laminate or a single ply scenario. Because the data files generated by RIPITx are generally numerous and large, the import process requires considerable computation time. However, once the data is read in, it is saved internally as part of the CAPP project file, eliminating the need for repeated imports during subsequent sessions.

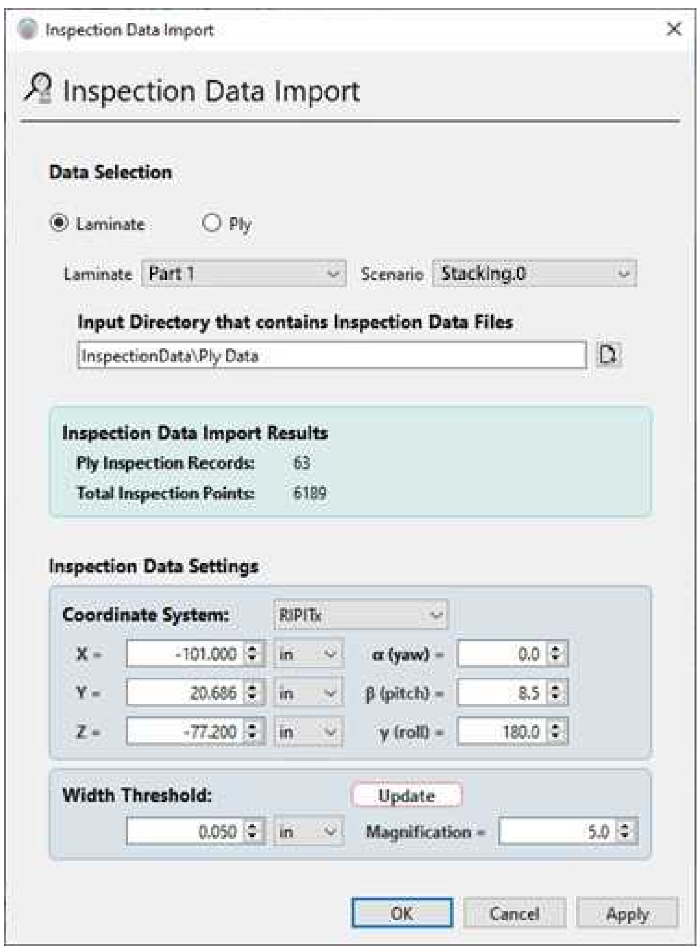

Figure 3.8: Inspection data import dialog within CAPP.

3.6.2.1 COORDINATE SYSTEM ALIGNMENT

A critical consideration during import is the coordinate system alignment between the tool surface model and the RIPITx frame of reference. The inspection data is collected in the machine coordinate system of the AFP cell, whereas the laminate model within CAPP resides in a part or tool surface coordinate system. Misalignment between these coordinate frames can lead to apparent spatial offsets or rotations between measured and predicted defect locations, misrepresenting the physical correlation between the two data sources.

To address this, CAPP allows the user to define a custom coordinate transformation that realigns the measurement data with the laminate reference frame. The transformation consists of a translation vector **t**, which specifies the origin offset of the imported inspection data, and three sequential Euler rotation angles that describe the orientation of the inspection coordinate frame relative to the laminate frame.

3.6.2.2 ROTATION MATRIX FORMULATION

The rotational component of the transformation is formulated using a Tait-Bryan Z-Y-X sequence expressed as a passive (frame) rotation.[3] Let α, β, and γ denote the Euler angles associated with rotations about the z, y, and x axes, respectively. The rotation matrix R is defined as the transpose of the active intrinsic Z-Y-X rotation:

$$R = \left(R_z(\alpha)\, R_y(\beta)\, R_x(\gamma)\right)^T \tag{3.75}$$

[3] In an active rotation, a vector is physically rotated within a fixed coordinate frame, whereas in a passive rotation the coordinate frame itself is rotated and the vector components are re-expressed in the new basis. The passive formulation is used here because the transformation maps inspection data expressed in a sensor-fixed coordinate system into the laminate reference frame, which is most naturally interpreted as a change of basis rather than a physical rotation of the geometry.

Expanding Equation (3.75) yields the explicit matrix form:

$$R = \begin{bmatrix} \cos\alpha\cos\beta & \sin\alpha\cos\beta & -\sin\beta \\ \cos\alpha\sin\beta\sin\gamma - \sin\alpha\cos\gamma & \sin\alpha\sin\beta\sin\gamma + \cos\alpha\cos\gamma & \cos\beta\sin\gamma \\ \cos\alpha\sin\beta\cos\gamma + \sin\alpha\sin\gamma & \sin\alpha\sin\beta\cos\gamma - \cos\alpha\sin\gamma & \cos\beta\cos\gamma \end{bmatrix} \quad (3.76)$$

3.6.2.3 AFFINE TRANSFORMATION

Let $\mathbf{p}_{\text{local}} = [x,\ y,\ z]^T$ denote a point expressed in the inspection coordinate system. The corresponding point in the laminate reference frame is obtained through the affine transformation:

$$\mathbf{p}_{\text{global}} = R\,\mathbf{p}_{\text{local}} + \mathbf{t} \quad (3.77)$$

where $\mathbf{t}$ is the translation vector defining the inspection frame origin relative to the laminate frame. This formulation ensures a physically meaningful mapping between coordinate systems. When alternate transformations are required, such as when processing inspection data acquired under different machine head orientations or tool configurations, the same formulation is applied with distinct rotation and translation parameters.

3.6.2.4 DEFECT ORIENTATION AND BINORMAL VECTOR

Each inspection point includes orientation data in the form of Euler angles that define the binormal direction of the detected defect. This orientation determines the direction along which the defect width is measured and displayed. Let ζ, ϕ, and θ denote the Euler angles provided by RIPITx for a given inspection point, corresponding to rotations about the x, y, and z axes of the sensor frame, respectively. The binormal vector $\mathbf{b}$ is computed as:

$$
\mathbf{b} = \begin{bmatrix} -\sin\theta\cos\phi \\ \cos\theta\cos\zeta - \sin\theta\sin\phi\sin\zeta \\ \cos\theta\sin\zeta + \sin\theta\sin\phi\cos\zeta \end{bmatrix} \tag{3.78}
$$

The binormal vector is used to generate the visual representation of each defect as a line segment centered at the measurement location, extending in both directions along **b** by half the measured width. When a coordinate transformation is applied to the inspection data, the binormal angles are adjusted accordingly to maintain consistent orientation in the laminate reference frame.

3.6.3 VISUALIZATION AND COMPARISON

After the RIPITx inspection data has been imported and processed, CAPP provides an integrated visualization environment that allows users to view both predicted and observed defects on the laminate surface simultaneously. This capability enables engineers to directly compare as-manufactured defects from RIPITx with as-designed predictions from VCP, providing insight into the quality of the layup and serving as an investigative tool when unforeseen defects occur.

Visualization of the inspection data utilizes colored line segments for the observed gaps and overlaps, where yellow represents detected gaps and purple represents detected overlaps. The location and orientation of each line segment are determined by the spatial coordinates and binormal vector computed from the RIPITx data, and the line lengths represent the reported defect widths. Predicted features from VCP are displayed concurrently, with green and red polygons outlining the anticipated gaps and overlaps, respectively. Dashed lines trace the course centerlines generated by VCP. The combination of these layers provides a direct comparison of as-designed and as-manufactured defect distributions.

Because RIPITx is configured to detect defects as small as 0.004 inches in width during the laydown process, the raw data can contain thousands of measurement points per ply. Two visualization parameters are provided to improve clarity and focus attention on significant defects. A width threshold can be defined to filter out defects below a specified size, excluding measurement points that fall below the threshold value. A magnification factor can also be applied to scale the displayed defect widths, as small features are often difficult to discern when projected onto a full-scale part model. These parameters can be adjusted interactively without modifying the underlying imported data, and the visualization updates in real time to reflect the current settings. Application of the visualization capability to the Quartic model, including discussion of the complementary detection scopes and process refinement implications, is presented in Section 4.6.

3.7 THE UNIFIED DIGITAL THREAD FOR AFP

The preceding sections have presented the individual components of the iSPP framework: subprocess integration with VCP for defect prediction and course generation (Section 3.1), the MIDS and DIDS formulations for quantifying manufacturing and structural quality (Sections 3.2 and 3.3), ply-level and laminate-level optimization algorithms for process improvement (Sections 3.4 and 3.5), and incorporation of inspection data for as-manufactured comparison (Section 3.6). Each component addresses a specific challenge in AFP process planning, yet the full value of the framework emerges from their integration into a cohesive system.

This section synthesizes these components into a unified digital thread that connects design intent, process planning, manufacturing execution, and quality verification. The term "digital thread" refers to the continuous flow of data and decision-making across traditionally separate disciplines, enabling traceability from initial design requirements through final part acceptance [44]. By establishing CAPP as the central integration platform, the iSPP framework transforms isolated tools into an interoperable system capable of closed-loop optimization and feedback.

3.7.1 PRIOR STATE: COMPLEMENTARY TOOLS AND EXISTING INTEGRATIONS

The software tools employed in the iSPP framework each represent mature, well-validated solutions for their respective domains. VCP provides industry-standard capabilities for AFP path generation, defect prediction, and NC code output. HyperX delivers aerospace-grade structural analysis and laminate optimization informed by manufacturing considerations. RIPITx offers high-resolution, in-situ inspection with automated defect detection. Each tool excels at its intended function and has been validated through extensive industrial application.

Prior to this work, integrations between these tools existed to address specific workflows. Collier Aerospace developed an interface between HyperX and VCP that maps as-manufactured fiber directions and lap/gap geometry from VCP to the finite element mesh in HyperX, enabling stress analysis that accounts for manufacturing deviations [196]. This integration supports iteration between structural sizing and path planning by translating manufacturing data into a form suitable for strength and stiffness evaluation. Similarly, Vericut developed a plugin for VCP that imports RIPITx profilometer data directly into the path programming environment, allowing visualization of inspection points overlaid on programmed courses to facilitate troubleshooting of manufacturing issues.

These existing integrations serve their intended purposes effectively. However, they represent point-to-point connections between specific tool pairs rather than a unified platform. Several capabilities remain unaddressed by these bilateral integrations:

- No unified metric exists to quantify process quality across manufacturing and structural domains simultaneously.
- No optimization algorithms operate on the integrated data to automatically improve

process plans.

- No single platform aggregates data from all sources (design, manufacturing simulation, and inspection) for comprehensive analysis.
- No closed-loop feedback connects inspection results back to process planning for iterative refinement.

The iSPP framework addresses these gaps by establishing CAPP as a central integration hub that connects the existing tools while adding scoring, optimization, and feedback capabilities that did not previously exist.

3.7.2 THE iSPP FRAMEWORK ARCHITECTURE

The iSPP framework employs a hub-and-spoke architecture with CAPP serving as the central integration platform. Rather than replacing the specialized capabilities of existing tools, CAPP orchestrates data flow between them while providing additional functionality for scoring, optimization, and visualization. Figure 3.9 illustrates the complete framework architecture, showing the relationships between external tools, internal modules, and the data flows that connect them.

The external tools—HyperX for structural analysis, VCP for path generation, and RIPITx for inspection—each interface with CAPP as described in Sections 3.1, 3.3, and 3.6, respectively. Within CAPP, the scoring modules (MIDS, DIDS, IDS) quantify process quality using the formulations presented in Sections 3.2 and 3.3, while the optimization modules (PLO, LLO) search the design space as described in Sections 3.4 and 3.5. This section focuses on how these components interact rather than reiterating their individual functions.

Table 3.5 summarizes the primary data exchanges within the iSPP framework, indicating the source, destination, and content of each interface.

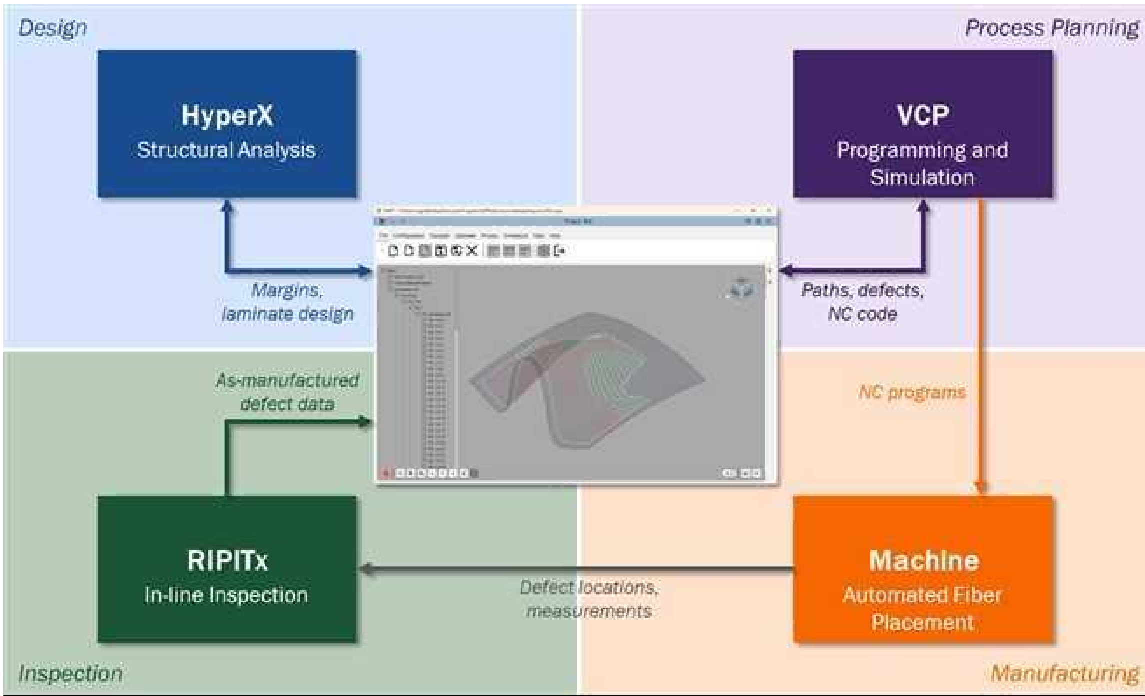

Figure 3.9: Architecture of the iSPP framework showing the integration of design, process planning, manufacturing, and inspection through the central CAPP platform.

Table 3.5: Summary of data flows within the iSPP framework.

From	**To**	**Data Exchanged**	**Section**
HyperX	CAPP	FE mesh, MoS values	3.3
CAPP	HyperX	Optimized scenario	3.3
CAPP	VCP	Ply scenarios	3.1
VCP	CAPP	Defect predictions, courses	3.1
VCP	AFP Machine	NC programs	3.1
AFP Machine	RIPITx	Manufactured ply	[59][a]
RIPITx	CAPP	Inspection data	3.6

[a]The citation refers to external documentation rather than a section of this dissertation.

The closed-loop nature of the framework is evident in two feedback paths. First, the CAPP-VCP loop enables iterative optimization: CAPP proposes ply scenario parameters, VCP evaluates them through path generation and defect prediction, and CAPP scores the results to guide subsequent iterations. This loop executes automatically during PLO and LLO, with the subprocess integration (Section 3.1.6) providing the computational efficiency required for high-throughput scenario evaluation. Second, the inspection feedback path allows comparison of as-manufactured defects with as-designed predictions, informing process refinement for future production or identifying systematic deviations that warrant design modification.

The framework's value lies in connecting these data flows through a unified platform. Without CAPP, the individual exchanges would remain isolated: structural margins would inform design but not process planning, defect predictions would guide programming but not optimization, and inspection data would identify problems but not feed back into systematic improvement. The iSPP architecture transforms these point-to-point connections into a coherent system where data flows continuously from design intent through manufacturing execution and part inspection.

3.7.3 MODULARITY AND EXTENSIBILITY

A deliberate design objective of the iSPP framework is modularity: each component interfaces with CAPP through defined communication patterns, allowing individual modules to be replaced or upgraded without disrupting the overall system. This plug-and-play architecture provides flexibility for both current use and future development.

3.7.3.1 INTERFACE ARCHITECTURE

The interfaces between CAPP and external tools are structured to localize tool-specific implementation details. The VCP integration communicates through VCP's Python API, with all API calls encapsulated within dedicated interface modules in CAPP. As described in Section 3.1.6, this architecture separates the API-specific commands from the higher-level optimization logic. If an alternative path generation tool were to replace VCP, only the interface module containing the API calls would require modification; the scoring algorithms, optimization routines, and visualization components would remain unchanged.

The HyperX integration follows a similar pattern, exchanging mesh geometry and margin data through .xml files that conform to a defined schema. The RIPITx integration processes tabular data files containing point-based defect measurements. While the current implementation reads TSV format, the parsing logic could accommodate other delimited formats (such as Comma-Separated Values (CSV)) with minimal modification, as the underlying data structure—point coordinates, defect type, and width—remains consistent across inspection systems.

3.7.3.2 REPLACEABLE COMPONENTS

The modular architecture supports replacement of individual components as improved solutions become available:

Alternative Path Generation Tools While VCP currently serves as the path programming engine, other AFP simulation software producing comparable defect predictions and course geometry could be integrated by developing a corresponding interface module that translates between the alternative tool's API and CAPP's internal data structures.

Alternative Structural Analysis Tools The HyperX integration could be adapted to other structural analysis platforms capable of providing element-level margin data on a compatible FEM, provided the alternative tool exports data in a format that can be mapped to the shared mesh discretization.

Alternative Inspection Systems The inspection data interface could accommodate other profilometer or vision systems that provide point-based defect measurements, with adjustments to the file parsing logic as needed.

Enhanced Optimization Algorithms The PLO and LLO algorithms represent initial implementations that demonstrate the framework's optimization capability. More sophisticated algorithms incorporating gradient-based methods, surrogate models, or multi-objective formulations could replace these modules while utilizing the same scoring metrics and VCP interface.

This modularity ensures that improvements to individual components—whether driven by advances in commercial software, new inspection technologies, or enhanced optimization methods—can be incorporated without requiring fundamental changes to the framework architecture.

3.7.4 SUMMARY

The iSPP framework unifies the individual components presented in this chapter into a coherent system for AFP process planning. By establishing CAPP as the central integration platform, the framework enables capabilities that were not achievable through isolated tools or point-to-point integrations: unified scoring metrics that translate between manufacturing and structural domains, automated optimization that iterates across the design space without manual intervention, and closed-loop feedback that connects inspection results back to process planning.

The framework builds upon the proven capabilities of existing tools while adding the integration layer and analytical modules necessary for smart process planning. The modular architecture ensures that improvements to individual components can be incorporated without disrupting the overall system, providing a foundation for continued development as AFP technology and computational methods advance. Chapter 4 verifies and demonstrates the complete framework through systematic application to the Quartic benchmark surface, exercising each component in sequence and demonstrating the integrated workflow from design intent through defect scoring, process optimization, and inspection visualization.

CHAPTER 4

CASE STUDY: THE QUARTIC BENCHMARK SURFACE

Chapter 3 presented the iSPP framework and the mathematical formulations underlying each of its components: the CAPP–VCP subprocess integration, the MIDS and DIDS scoring algorithms, the PLO and LLO optimization routines, and the inspection data visualization capability. The present chapter verifies and demonstrates these components in sequence through application to a single benchmark geometry, the Quartic model, with the objective of confirming functional correctness, quantifying performance characteristics, and illustrating the integrated workflow that connects design through process planning.

The chapter begins with a description of the Quartic benchmark geometry and the laminate configuration used throughout the case study (Section 4.1). The verification and demonstration is then organized in three levels. At the component level, data exchange fidelity is confirmed by comparing outputs from the CAPP–VCP subprocess integration against those produced by native VCP, and by verifying that structural margin data imported from HyperX is preserved without loss or distortion (Section 4.2). At the scoring level, the discriminative capability of the MIDS formulation is demonstrated by comparing two distinct ply scenarios, and the effect of margin-weighted scoring on the DIDS and IDS is shown through a complete structural feedback loop (Section 4.3). At the optimization level, the PLO algorithm is exercised on representative plies to characterize the design space and the algorithm's behavior (Section 4.4), and both LLO severity kernels—the stagger shift kernel and the gravity-based kernel—are applied to the full 11-ply laminate to demonstrate defect redistribution across the stacking sequence and to evaluate the comparative behavior of the two interaction models (Section 4.5). The chapter concludes with a demonstration of the inspection data visualization capability (Section 4.6) and an integrated discussion of the results across all components (Section 4.7).

4.1 BENCHMARK GEOMETRY: THE QUARTIC SURFACE

The Quartic model is a doubly-curved saddle-shaped surface named for its quartic polynomial definition. Collier Aerospace developed this geometry as a benchmark for AFP process planning, with aggressive curvature designed to generate significant fiber deviations and gap/overlap defects during path generation. The model spans approximately 100 in. by 100 in. in plan view, with curvature that produces challenging manufacturing conditions representative of complex aerospace structures. Figure 4.1 presents the CAD geometry and corresponding FEM.

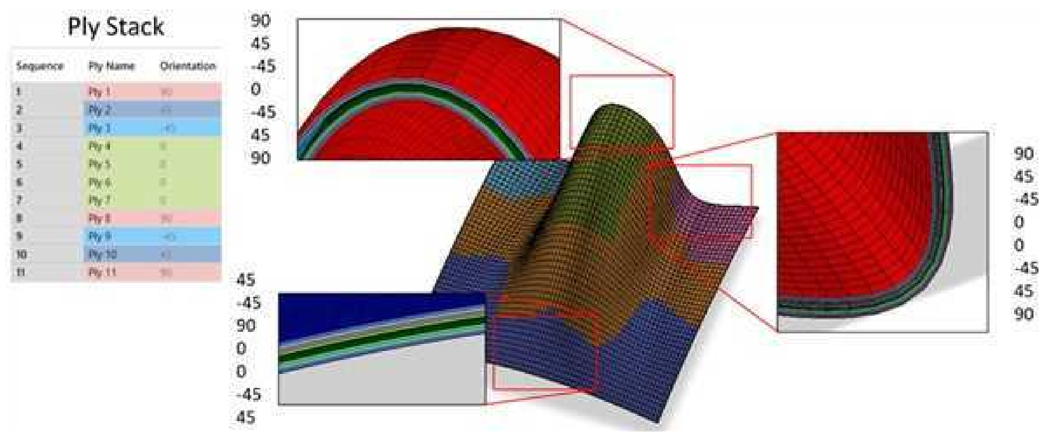

Figure 4.1: CAD geometry and finite element mesh for the Quartic model, provided by Collier Aerospace.

The laminate configuration consists of an 11-ply layup. Four distinct ply boundary types define the spatial extent of individual plies: full-coverage plies that span the entire surface, partial plies covering a subset region, partial complement plies covering the remaining area, and corner plies covering localized sub-regions of the surface. Table 4.1 summarizes the stacking sequence with corresponding ply boundaries and baseline starting points. Figure 4.2 illustrates the four boundary types. The baseline ply scenario starting points listed in Table 4.1 were provided by Collier Aerospace as part of the benchmark definition and represent state-of-the-art manual process planning practice: each starting point was selected by an experienced engineer using domain knowledge and judgment within VCP, without the benefit of automated optimization.

Table 4.1: Stacking sequence and ply boundary definitions for the Quartic model.

Ply	**Fiber Angle**	**Ply Boundary**	**Baseline Starting Point (in.)**
1	90°	Partial	$(-3.977, -9.155, 33.186)$
2	45°	Full	$(5.057, 0.879, 28.096)$
3	−45°	Full	$(5.057, 0.879, 28.096)$
4	0°	Full	$(5.057, 0.879, 28.096)$
5	0°	Bottom left corner	$(36.173, -35.101, -3.064)$
6	0°	Top left corner	$(-34.489, -33.115, -3.419)$
7	0°	Partial complement	$(-0.593, 35.021, 13.367)$
8	90°	Partial complement	$(-0.593, 35.021, 13.367)$
9	−45°	Full	$(5.057, 0.879, 28.096)$
10	45°	Full	$(5.057, 0.879, 28.096)$
11	90°	Partial	$(-3.977, -9.155, 33.186)$

Process planning for the Quartic model employs eight 0.5 in. tows per course, yielding a 4 in. course width. The FEM comprises 3,645 quadrilateral elements that serve as the common spatial reference for defect mapping, structural margin assignment, and score computation. Combined in-plane force resultants and bending moment resultants, provided by Collier Aerospace, are applied for structural analysis in HyperX, providing element-level MoS values that inform the margin-weighted scoring demonstrated in Section 4.3. The applied load resultant distributions are shown in Figure 4.3. Because the Quartic model is a virtual benchmark and was not physically fabricated, the inspection demonstration in Section 4.6 employs synthetic data to exercise the visualization and import capabilities.

The ply boundary structure of the Quartic laminate has a direct bearing on the LLO results presented in Section 4.5. The partial boundary of Ply 1, together with the partial complement boundary of Ply 8, constitutes the full laminate footprint in the 90° orientation. Plies 2 and 10 (+45°) and Plies 3 and 9 (−45°) each consist of two full-boundary plies whose

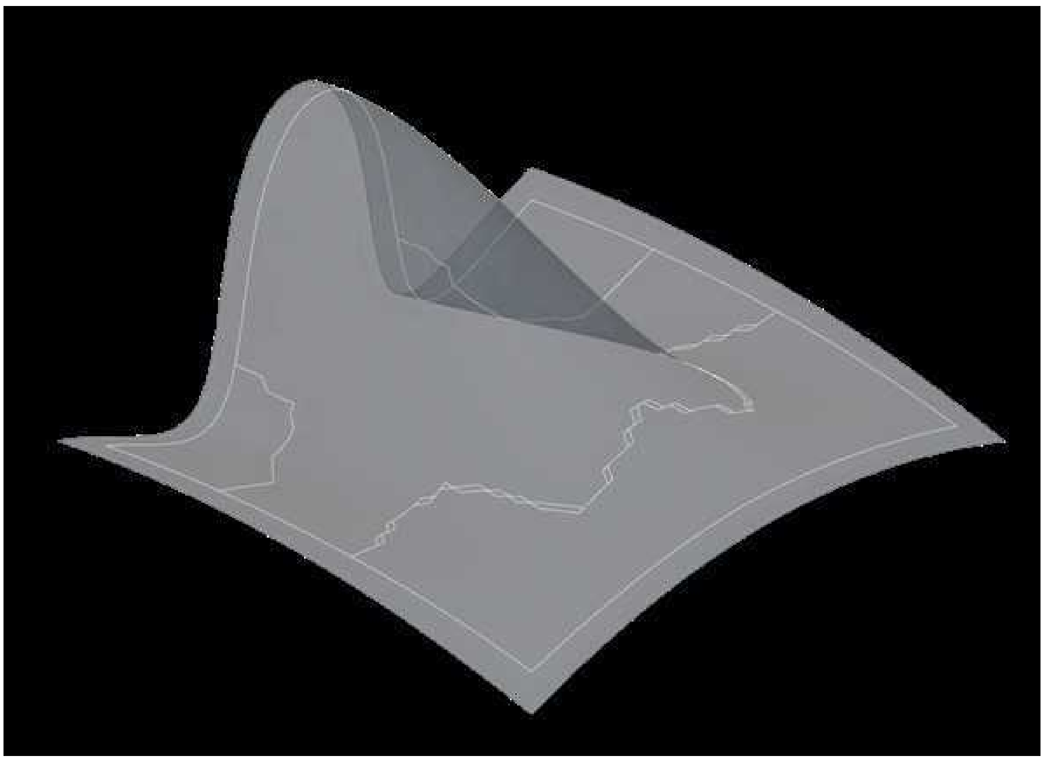

Figure 4.2: Ply boundary definitions for the Quartic model: full coverage, partial, partial complement, and corner regions.

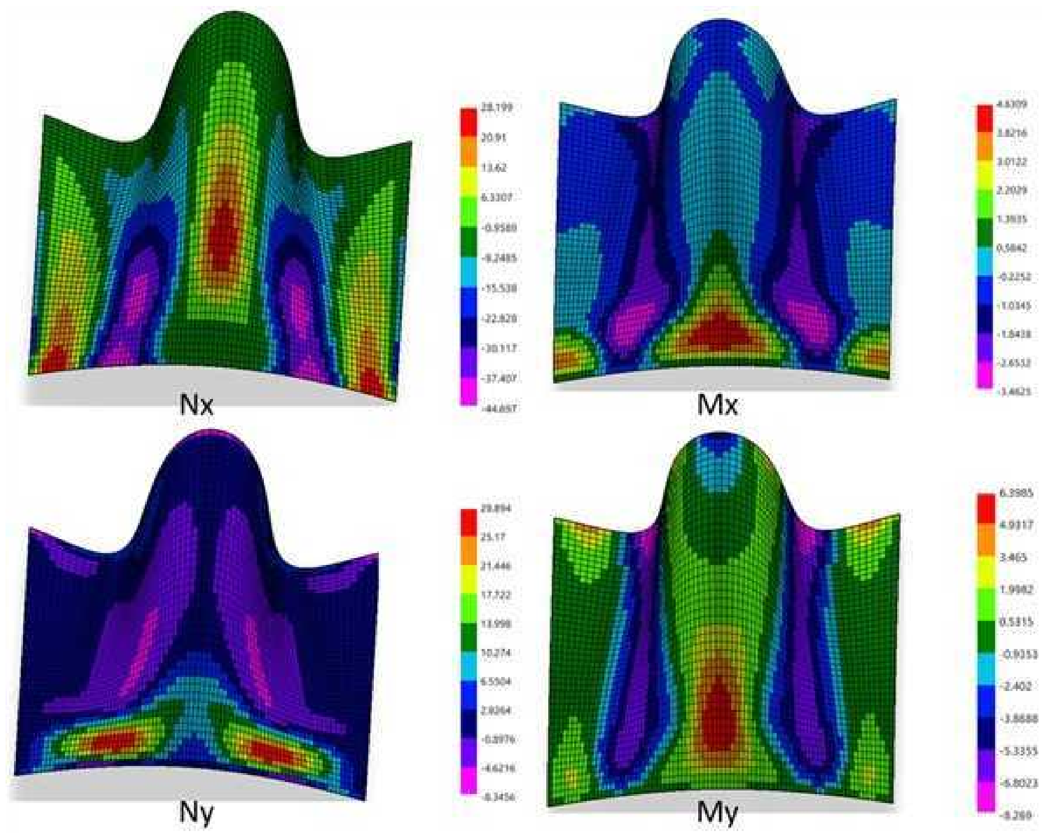

Figure 4.3: Applied load resultant distributions for the Quartic model: in-plane force resultants N_x and N_y and bending moment resultants M_x and M_y, provided by Collier Aerospace.

defect distributions are directly overlaid in the baseline configuration. The four 0° plies (Plies 4 through 7) span boundaries of different extents, with Ply 4 covering the full surface area, Plies 5 and 6 occupying the bottom-left and top-left corner sub-regions respectively, and Ply 7 covering the partial complement region. These structural relationships—particularly the co-located full-boundary ply pairs at ±45°—represent the primary defect stacking risk in the baseline laminate and the primary target for LLO improvement.

4.2 CAPP–VCP INTEGRATION VERIFICATION

The subprocess-based CAPP–VCP integration described in Section 3.1.6 introduces a headless execution pathway that bypasses VCP's graphical interface. Before this integration can be trusted as the simulation kernel for scoring and optimization, it must be shown to produce outputs that are functionally identical to those generated by native VCP operating under direct user control. Any discrepancy would introduce a systematic bias into all downstream scoring and optimization computations.

4.2.1 VERIFICATION METHOD

The verification study used Quartic model Ply 1 as the test case, employing a single, fixed set of inputs including the surface geometry, composite process definition file, and the ply scenario inputs for Ply 1 as defined in Table 4.1. Two `.VCLaminate` files were generated independently: one through the native VCP environment operated manually, and one through CAPP's headless subprocess call. The resulting files were compared across three categories: course generation outputs, defect prediction outputs, and ply-level scoring outputs.

4.2.2 VERIFICATION RESULTS

Table 4.2 presents the comparison results for Ply 1. The two integration methods produced agreement across every reported metric. Course geometry attributes—including tow count, total tow length, headpath length, headpath endpoints, and service cut count—were identical. Defect occurrence counts for all four defect types (gaps, overlaps, angle deviation exceedances, and steering violations) matched precisely. Ply-level scores computed from the imported defect data agreed to two decimal places, the full precision reported by CAPP's scoring module, confirming that no scoring discrepancy was introduced by the integration pathway.

Table 4.2: Verification of CAPP–VCP subprocess integration using the Quartic model (Ply 1).

Metric	Native VCP[a]	CAPP–VCP Integration
Course generation outputs		
Course details	8 tows, 1 packet	8 tows, 1 packet
Tow width (in.)	0.50	0.50
Total tow length (yd.)	34.33	34.33
Head-path length (yd.)	4.29	4.29
Head-path start point (in.)	(45.07, −31.87, 1.41)	(45.07, −31.87, 1.41)
Head-path end point (in.)	(−45.07, −31.87, 1.41)	(−45.07, −31.87, 1.41)
Service cuts (count)	16	16
Defect prediction outputs		
Gap occurrences	32	32
Overlap occurrences	51	51
Angle deviation occurrences	2,511	2,511
Steering occurrences	2,472	2,472
Scoring outputs		
Total ply score	98.51	98.51
Manufacturing score	98.45	98.45
Margin-of-safety score	98.55	98.55

[a]Laminate file generated by standalone VCP, then loaded into CAPP for scoring.

Figure 4.4 provides a qualitative comparison through side-by-side visualization of the headpath geometry and defect maps generated by each method. The course centerlines, gap polygon locations, and overlap polygon locations are spatially coincident across both outputs, confirming that the subprocess architecture does not introduce any geometric perturbation.

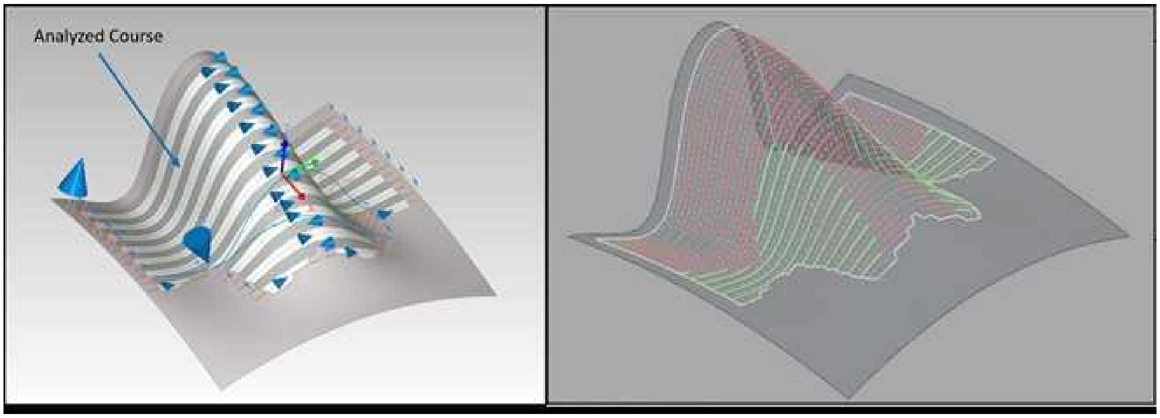

Figure 4.4: Qualitative verification of the CAPP–VCP subprocess integration for Quartic model Ply 1, showing matching outputs between native and integrated workflows.

4.2.3 SUBPROCESS EFFICIENCY

Beyond output fidelity, the subprocess architecture provides a substantial efficiency advantage over the socket-based integration used in prior work [203]. The socket-based method required complete laminate regeneration for every parameter change, because it created a new laminate file and redefined all process parameters before regenerating every ply serially. The subprocess architecture maintains a persistent VCP process in memory, allowing CAPP to modify and regenerate individual plies while leaving all others untouched.

The efficiency improvement was quantified by processing 63 ply scenarios for Quartic model Ply 1 under both methods. The socket-based approach (Method 1) required 4.54 hours due to serial reinitialization of the complete laminate at each iteration. The subprocess-based method (Method 2) completed the same 63 scenarios in 23.65 minutes—a reduction of 91.3% in total processing time (Figure 4.5).[1] This order-of-magnitude reduction transforms the computational budget available for optimization: scenarios that would have required multi-day computation under the prior architecture are completed within a single working session.

[1]Computation times were obtained on a laptop equipped with an Intel® Core™ Ultra 7 155H processor and 32 GB RAM.

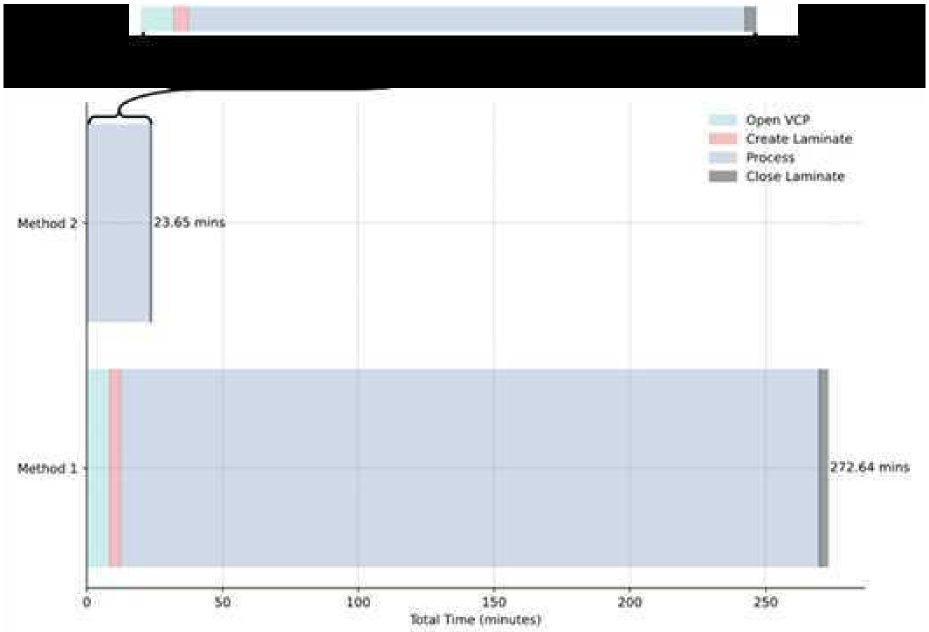

Figure 4.5: Comparison of processing time between socket-based (Method 1) and subprocess-based (Method 2) integration architectures for 63 iterations of Quartic model Ply 1.

The combination of exact output fidelity and 91.3% efficiency improvement establishes the subprocess integration as the verified computational foundation for all subsequent scoring and optimization studies in this chapter. All VCP results reported in the following sections were generated through the subprocess interface.

4.3 DEFECT SCORING VERIFICATION

With the CAPP–VCP integration verified, this section demonstrates and verifies the scoring formulations developed in Sections 3.2 and 3.3. The verification proceeds in three stages: the MIDS formulation is verified by demonstrating its discriminative capability across two distinct ply scenarios; the HyperX–CAPP data exchange is verified by confirming the integrity of imported structural margin data; and the DIDS and IDS formulations are demonstrated through a complete structural feedback loop that closes from defect-aware optimization back to structural reanalysis.

4.3.1 MIDS: SCORING DISCRIMINABILITY

A score that cannot distinguish between better and worse process plans has no value as an optimization objective. The MIDS verification therefore focuses on demonstrating that the formulation responds correctly to changes in defect distribution between two candidate ply scenarios. The equal-weight defect criteria configuration from Table 3.3 was used throughout, with thresholds of 0.125 in. for gaps and overlaps, 5.0° for angle deviation, and 15.0 in. minimum steering radius.

4.3.1.1 SCENARIO DEFINITIONS

Two ply scenarios were evaluated for Quartic model Ply 1. Scenario A uses the baseline starting point listed in Table 4.1. Scenario B uses an alternative starting point intentionally selected as a poor-performing candidate to demonstrate the discriminative capability of the MIDS formulation.

4.3.1.2 RESULTS

Table 4.3 presents the defect statistics and computed MIDS components for both scenarios. Figure 4.6 shows the corresponding defect visualizations.

The MIDS values of 98.45 and 95.72 for Scenarios A and B, respectively, demonstrate that the formulation successfully discriminates between ply scenarios with different defect characteristics. The 2.73-point differential arises from the increased gap and overlap area in Scenario B, which produces higher frequency and severity penalties in those defect type components relative to Scenario A. The component penalty decomposition provides diagnostic information identifying the specific defect types responsible for the score difference. Both scenarios were evaluated using the rosette layup strategy, which is discussed further in the context of the PLO results in Section 4.4.

A natural observation from these results is that both scenarios score high, and the

Table 4.3: Comparison of MIDS computation for two Quartic model Ply 1 scenarios.

Metric	Scenario A	Scenario B
Defect counts (exceeding threshold)		
Gap occurrences	32 (30)	36 (34)
Overlap occurrences	51 (41)	63 (44)
Angle deviation occurrences	2,511 (0)	2,456 (0)
Steering occurrences	2,472 (0)	2,378 (0)
Frequency penalties		
$P_{\text{gap}}^{\text{freq}}$	0.012	0.052
$P_{\text{overlap}}^{\text{freq}}$	0.018	0.057
$P_{\text{angle}}^{\text{freq}}$	0.000	0.000
$P_{\text{steering}}^{\text{freq}}$	0.000	0.000
Severity penalties		
$P_{\text{gap}}^{\text{sev}}$	0.049	0.116
$P_{\text{overlap}}^{\text{sev}}$	0.045	0.117
$P_{\text{angle}}^{\text{sev}}$	0.000	0.000
$P_{\text{steering}}^{\text{sev}}$	0.000	0.000
Aggregated score		
MIDS	98.45	95.72

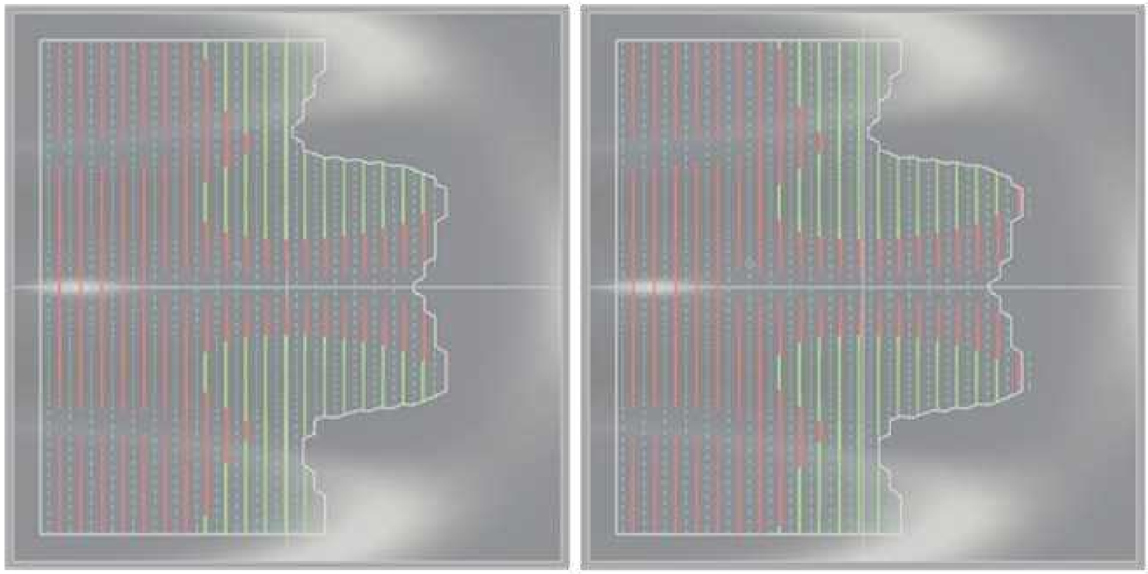

Figure 4.6: Defect visualization for Quartic model Ply 1: (a) Scenario A (baseline) and (b) Scenario B (alternative starting point). Green polygons indicate gaps; red polygons indicate overlaps.

difference between them—while directionally correct and physically meaningful—represents a relatively narrow numerical range. This behavior is not a deficiency of the metric. Because the Quartic model's gap and overlap defects arise primarily at course boundaries and are governed by the surface curvature rather than by the choice of starting point, all feasible ply scenarios cluster in the upper portion of the $[0, 100]$ score range. Starting-point variation redistributes defects spatially but does not alter the underlying defect-generating mechanism; the total defect area across the ply is approximately conserved, and score differences between scenarios reflect its redistribution rather than its creation or elimination. The practical implication is that score differentials of 2–8 points represent the full range of achievable improvement through starting-point optimization on this geometry, and the optimization algorithm's task is to reliably identify the upper end of this range. Section 4.4 demonstrates that the PLO algorithm does so consistently.

4.3.2 HYPERX-CAPP DATA EXCHANGE VERIFICATION

The DIDS formulation requires element-level MoS data from HyperX mapped onto the shared FEM. Any error in this data exchange—node count mismatches, unit scaling errors, or identifier mapping failures—would corrupt the margin-weighted penalty calculations. A verification study was therefore conducted to confirm that structural data exported from HyperX is imported into CAPP without loss or distortion.

The verification used the Quartic model Stacking.0 laminate scenario. The FEM, optimization settings, and MoS data were exported from HyperX and imported into CAPP. A record-by-record comparison of node coordinates, element connectivity, and MoS values confirmed complete data integrity.

Table 4.4: Verification of HyperX–CAPP data exchange for the Quartic model.

Attribute	**HyperX Export**	**CAPP Import**
Nodes	3,772	3,772
Elements	3,645	3,645
Units	in.	in.
Plies	11	11
MoS minimum	−0.31	−0.31
MoS maximum	52.86	52.86
Layup strategies	Rosette	Rosette
Defect criteria	Gap, Overlap, Angle Deviation, Steering	

Table 4.4 confirms exact agreement across all verified attributes. The MoS range (−0.31 to 52.86) is preserved without scaling or offset, confirming that the unit system and element identifier mapping are correctly handled. The presence of negative MoS values in the baseline configuration—86 elements have MoS < 0, primarily concentrated near the ply drops at the base of the saddle geometry—motivates the FoS transformation described in Section 3.3.2 and illustrates why margin-weighted scoring is a meaningful addition to pure manufacturability assessment.

4.3.3 DIDS AND IDS: MARGIN-WEIGHTED SCORING AND THE STRUCTURAL FEEDBACK LOOP

The DIDS and IDS formulations are verified through a three-stage demonstration: baseline characterization, IDS-optimized process planning, and structural reanalysis with local refinement. Together, these stages demonstrate the closed-loop interaction between CAPP and HyperX that constitutes the design-aware component of the iSPP framework.

4.3.3.1 BASELINE CONFIGURATION

The Collier Aerospace baseline process plans introduced in Section 4.1 serve as the baseline here; these plans were generated in VCP without margin-informed optimization and therefore reflect current manual practice prior to any IDS-driven refinement. The baseline MoS distribution, shown in Figure 4.7, reveals a pronounced concentration of negative margins along the left side of the saddle-shaped geometry, in the region of the partial ply drops. Tow gaps are concentrated across multiple regions of the surface, with notable density in the ±45° plies where fiber steering and surface curvature are most pronounced. The accumulated effect of these defects produces localized structural capacity reductions that drive the 86 negative-margin elements identified during data exchange verification. The global MoS range for the baseline configuration is −0.31 (minimum) to 52.86 (maximum).

4.3.3.2 IDS-OPTIMIZED CONFIGURATION

CAPP's IDS-based optimization was applied to the Quartic laminate, adjusting ply starting points and layup strategies to maximize the IDS weighted objective. The optimization operates through the PLO and LLO algorithms described in Sections 3.4 and 3.5, with $w_{\text{design}} = w_{\text{mfg}} = 0.5$ to balance manufacturing quality and structural sensitivity. The resulting MoS distribution is shown in Figure 4.8.

The concentration of tow gaps previously located near the ply-drop region on the

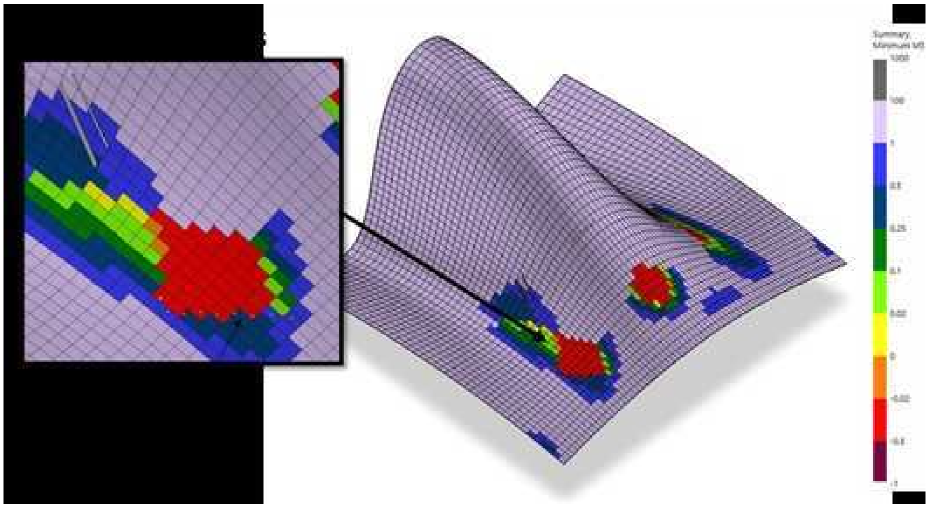

Figure 4.7: Baseline MoS distribution for the Quartic model. Negative margins (red) are concentrated at the ply-drop region on the left side of the saddle geometry.

left side of the model was significantly reduced, leading to improved structural margins in that area. The global MoS range remained at 52.86 (maximum) and −0.31 (minimum), reflecting that the absolute worst-case element was not altered by starting-point redistribution alone. However, the number of finite elements exhibiting negative margins decreased by approximately 30%, from 86 to 60 elements. This reduction reflects a more uniform redistribution of defects across the structural domain rather than their elimination: the optimization shifts defect concentrations away from the most structurally critical region, improving the overall margin landscape without requiring additional material.

4.3.3.3 STRUCTURAL REANALYSIS AND LOCAL REINFORCEMENT

The margin-weighted optimization successfully reduced the concentration of critical defects but did not eliminate all negative margins, because the residual negative margins in the optimized configuration are attributable to intrinsic geometric and loading conditions that starting-point variation cannot resolve. To close the design loop, the optimized laminate configuration and corresponding as-manufactured fiber directions were reimported into HyperX for full structural reanalysis.

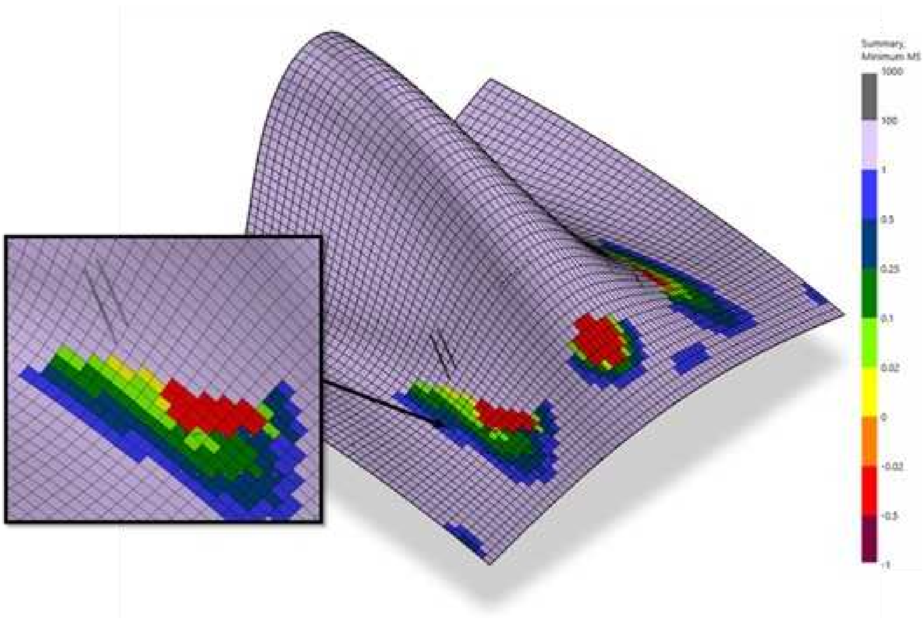

Figure 4.8: MoS distribution after IDS-based optimization, showing reduced concentration of negative margins at the ply-drop region.

The reanalysis confirmed that the redistributed layup improved overall compliance with allowable limits, but localized overstressing persisted at the base of the saddle geometry due to the interaction of surface curvature with steering-induced fiber deviation. To resolve these residual negative margins, HyperX's laminate optimization capability was used to introduce local reinforcement. Two 90° plies were added symmetrically on the second and twelfth layers of the laminate, targeting the low-margin region identified in the MoS map (Figure 4.9).

After insertion of the reinforcement plies, all margins across the structure became positive. The minimum MoS increased to 0.437, located at the left side of the saddle base, and the full margin distribution is shown in Figure 4.10.

4.3.3.4 SUMMARY OF FEEDBACK LOOP

The three-stage sequence—baseline characterization, IDS-optimized process planning, and structural reanalysis—demonstrates the collaborative interaction the iSPP frame-

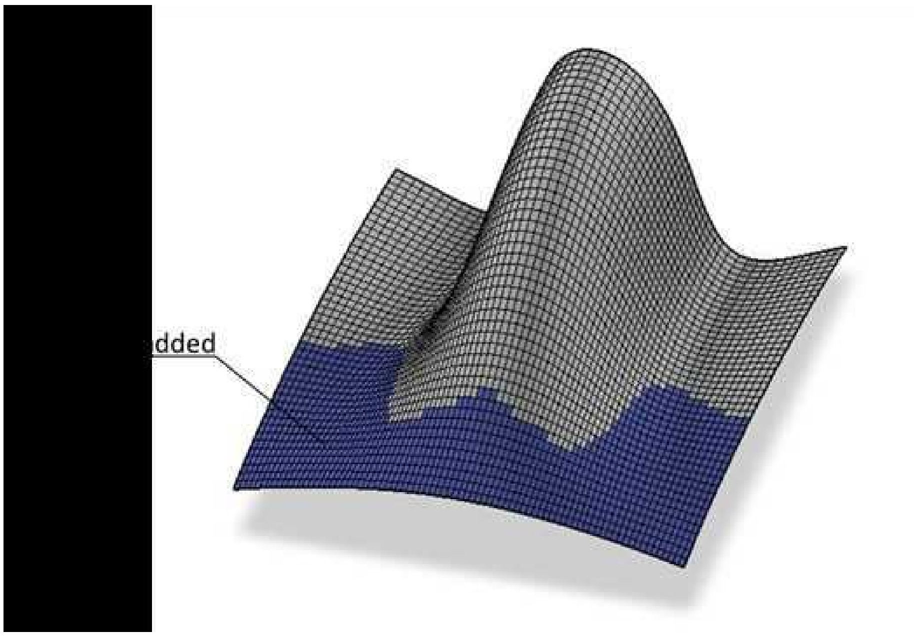

Figure 4.9: Regions of local reinforcement added to the Quartic model. Two 90° plies were inserted at layers 2 and 12 to resolve residual negative margins at the saddle base.

work enables. The CAPP–HyperX digital thread allows defect distributions predicted by VCP to inform structural analysis, structural margins to inform defect weighting, and optimized process plans to be verified and refined within the same workflow. The 30% reduction in negative-margin elements achieved through process plan optimization alone underscores the value of margin-weighted scoring: the improvement was obtained not by adding material but by redistributing existing defects away from structurally critical regions, a result that is only achievable when the optimization objective explicitly reflects where defects matter.

4.4 PLY-LEVEL OPTIMIZATION RESULTS

The PLO algorithm was applied to the Quartic model to evaluate its behavior on a surface with aggressive curvature, characterize the IDS landscape that drives the optimization, and demonstrate the algorithm's two-phase search strategy. Results are presented for Ply 1 (90°, partial boundary) and Ply 2 (45°, full boundary), which together represent the two principal ply boundary types in the Quartic laminate.

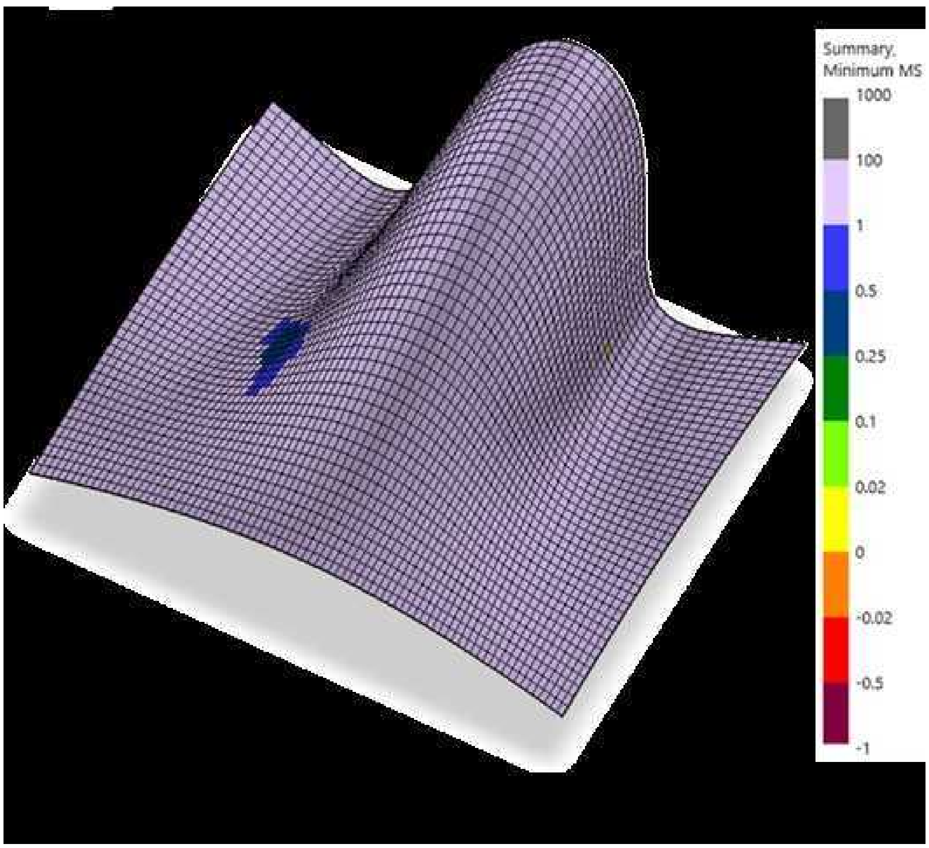

Figure 4.10: Final MoS distribution after local reinforcement, showing all-positive margins with minimum MoS = 0.437.

4.4.1 BASELINE CONFIGURATIONS

The baseline ply scenarios were generated by VCP using the starting points listed in Table 4.1. These baselines represent the as-designed configuration against which optimization improvements are measured. Figure 4.11 presents the defect visualizations and associated IDS values for both plies.

4.4.2 SURVEY PHASE

The rosette layup strategy was selected for the Quartic model throughout all PLO and LLO studies. The saddle-shaped surface lacks a single dominant flat region from which a fixed fiber reference direction would project consistently, and the rosette strategy's geodesic propagation from a user-defined origin adapts to the local surface curvature in a manner that uniform angular projection does not. A fixed-angle strategy on this geometry produces systematic fiber deviations in regions of high Gaussian curvature that register as angle deviation defects regardless of starting point; the rosette strategy minimizes this effect by construction and was therefore used as the sole layup strategy for all optimization studies reported in this chapter.

The global survey was executed with a grid spacing parameter of γ = 310% of the course width for this demonstration, increased from the production default of 150% to 310% to reduce computation time while still providing sufficient spatial coverage to illustrate the algorithm's behavior. At 310%, the grid produces approximately one survey point per two course widths, ensuring that the major regions of the ply domain are sampled. For Ply 1, this spacing produced 58 survey scenarios across the partial ply boundary; for Ply 2, the larger full boundary yielded 69 survey scenarios. Each grid point was evaluated with the rosette layup strategy, as this is the single strategy configured for the Quartic model. The distribution of survey starting points for both plies is shown in Figure 4.12.

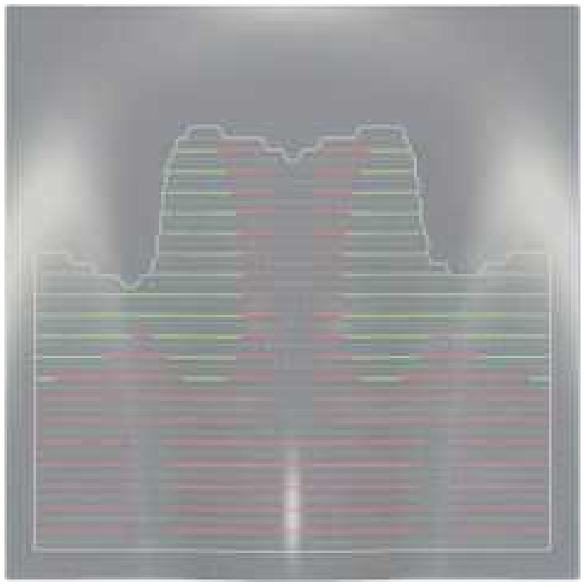

(a) Ply 1 (90°, partial boundary) — IDS: 98.51

(b) Ply 2 (45°, full boundary) — IDS: 97.94

Figure 4.11: PLO baseline defect distributions for the Quartic model. Green regions indicate gaps; red regions indicate overlaps.

4.4.3 OPTIMIZATION RESULTS

The IDS landscape characterization described in Section 3.4.6 confirmed that the perpendicular-shift profile is discontinuous for both ply orientations on the Quartic surface,

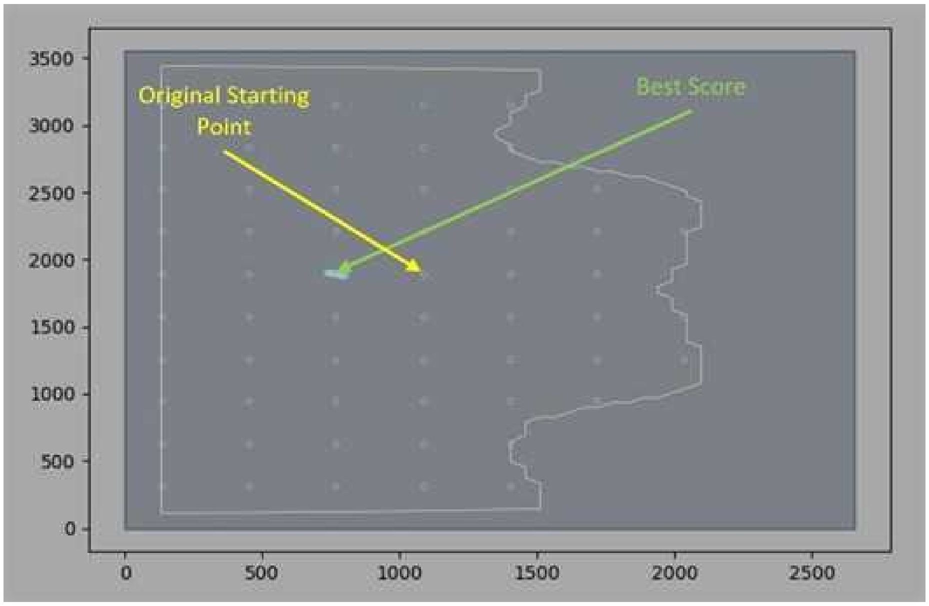

(a) Ply 1 (90°)

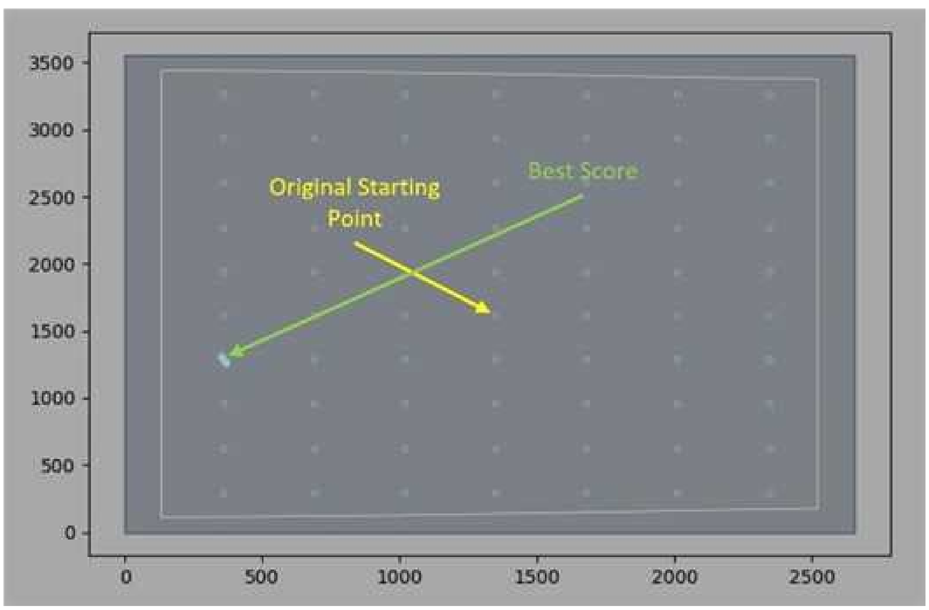

(b) Ply 2 (45°)

Figure 4.12: Distribution of survey starting points (grid) and perpendicular-shift candidates (cluster) for the Quartic model. The tightly spaced cluster indicates the local refinement region around the highest-scoring survey seed.

validating the sampling-based refinement strategy over interpolation-based approaches. The optimization results are as follows.

After the survey phase, the five highest-scoring spatially diverse seeds were selected for perpendicular-shift refinement using the spatial diversity criterion described in Section 3.4.4. The perpendicular-shift phase generated 24 additional candidates per seed (n = 25 shift points total, including the seed). Total scenario counts were 82 for Ply 1 and 93 for Ply 2, including both survey and shift candidates.

The optimized defect distributions are shown in Figure 4.13, and Table 4.5 provides a numerical comparison of baseline and optimized IDS values.[2]

Table 4.5: Baseline and PLO-optimized IDS comparison for the Quartic model.

Ply	**Angle**	**Boundary**	**Baseline IDS**	**Optimized IDS**	**Improvement**
1	90°	Partial	98.51	98.52	+0.01 (+0.010%)
2	45°	Full	97.94	98.05	+0.11 (+0.112%)

4.4.4 DISCUSSION

The optimized scenarios exhibit marginal IDS improvements over their respective baselines. For Ply 1, the improvement of 0.01 points approaches the precision limit of the scoring module and should be interpreted as the algorithm identifying a configuration at least as good as the baseline rather than a meaningful quality improvement. For Ply 2, the 0.11-point improvement represents a more substantive redistribution of defects within the available design space. Both results are consistent with the slight reductions in the number of gaps and overlaps exceeding the threshold value observed in the defect visualizations. This result is consistent with the analysis presented in Section 3.4.8: because both baseline and optimized scenarios use the same number of tows and the same layup strategy, the defect distribution is primarily governed by the tool surface geometry. Starting-point variation redistributes defects spatially but has limited capacity to reduce their total occurrence for a

[2] Computation times were approximately 33 minutes for Ply 1 and 38 minutes for Ply 2, obtained on a laptop equipped with an Intel® Core™ Ultra 7 155H processor and 32 GB RAM.

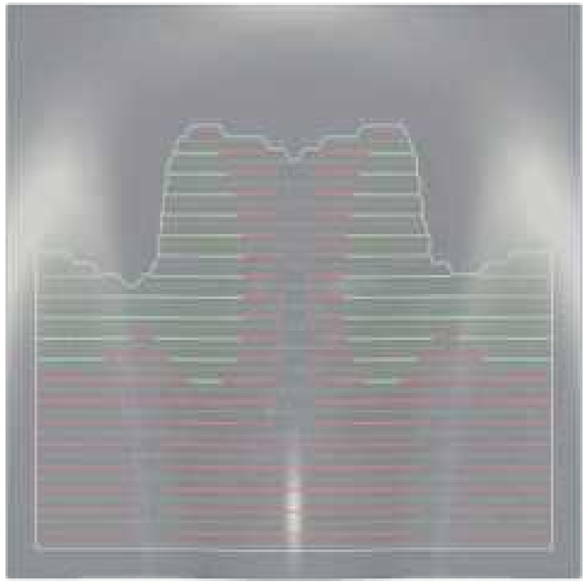

(a) Ply 1 (90°) — IDS: 98.52

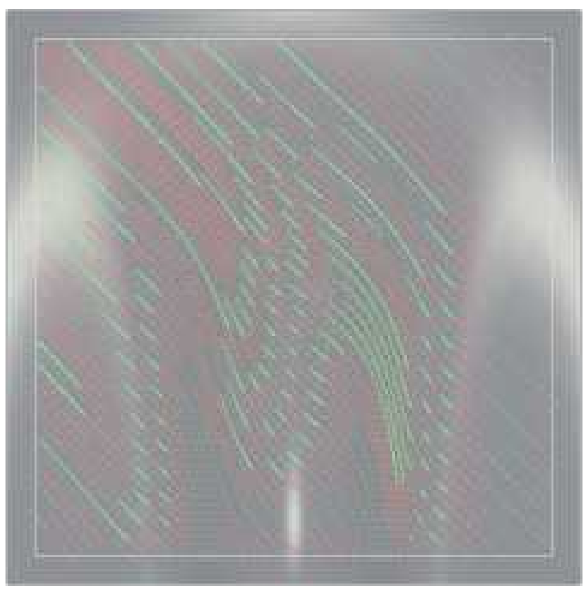

(b) Ply 2 (45°) — IDS: 98.05

Figure 4.13: PLO optimized defect distributions for the Quartic model.

single ply considered in isolation. This redistribution behavior is a necessary precondition for LLO: because the defect pattern shifts with starting point while the total defect area remains approximately conserved, different ply scenarios present structurally distinct through-thickness defect profiles. LLO exploits precisely this variation, selecting combinations of ply scenarios whose defect distributions are spatially offset across the laminate thickness, to achieve stagger configurations that minimize defect coincidence.

This observation connects directly to the score range discussion in Section 4.3.1. The narrow numerical range of IDS values across all survey scenarios is not a failure of the optimization—it reflects the physical constraint that gap and overlap defects at course boundaries are largely determined by the surface curvature rather than by the choice of starting point. The PLO algorithm correctly identifies the best achievable configuration within this constrained landscape.

The more consequential role of PLO in the Quartic validation is as a scenario generation engine. The survey and perpendicular-shift phases produce a diverse population of scored ply scenarios across the full extent of each ply's domain. This population constitutes the candidate pool that enables LLO to select ply combinations that minimize through-thickness defect stacking—an objective for which the marginal per-ply improvements documented here are less important than the spatial diversity of the candidate set. The LLO results presented in Section 4.5 demonstrate that this diversity translates into measurable laminate-level improvement.

4.5 LAMINATE-LEVEL OPTIMIZATION RESULTS

The LLO framework presented in Section 3.5 was applied to the complete 11-ply Quartic laminate under both severity kernel configurations to demonstrate the framework's ability to reduce through-thickness defect stacking and to evaluate the comparative behavior of the stagger shift and gravity-based kernels. Unlike PLO, which operates on individual plies in isolation, LLO evaluates all ply scenario combinations jointly and selects the laminate configuration that minimizes the spatial coincidence of defects across the stacking sequence. This section presents the shared baseline configuration and candidate pool, the optimized result under each kernel, and a head-to-head comparison of the two kernels with respect to score improvement, optimized ply assignments, and defect redistribution behavior.

The presentation proceeds as follows. Section 4.5.1 establishes the baseline laminate configuration and baseline L-IDS values under both kernels. Section 4.5.2 describes the candidate pool generation and, for the gravity-based kernel, the neighbor precomputation required prior to optimization. Sections 4.5.3 and 4.5.4 present the optimization results for the stagger shift and gravity-based kernels, respectively, including the optimized ply scenario assignments, score improvements, and defect stacking visualizations. Section 4.5.5 provides a direct quantitative and qualitative comparison of the two kernel outputs. Section 4.5.6 synthesizes the findings and discusses their implications for the selection and application of the two kernels in practice.

4.5.1 BASELINE LAMINATE CONFIGURATION

The baseline laminate scenario was constructed using the starting points defined in Table 4.1, consistent with the Collier Aerospace baseline introduced in Section 4.1. This baseline is common to both kernel evaluations and establishes the reference against which both optimization results are compared.

The baseline defect distribution across the full laminate is shown in Figure 4.14. Defect stacking is clearly visible as high-density regions where gaps and overlaps from multiple plies coincide spatially.

The stacking behavior becomes most apparent when the defect data is disaggregated by fiber orientation, as shown in Figure 4.15. For orientations other than 0°, the baseline definition assigns identical starting points to all plies within each angle group, so the resulting defect polygons are directly overlaid: the +45° defects of Ply 2 stack precisely on those of Ply 10, and the −45° defects of Ply 3 stack on those of Ply 9. This direct through-thickness coincidence produces the maximum possible stacking severity for these ply pairs under both kernels and represents the worst-case configuration for the LLO objective. The 0° plies,

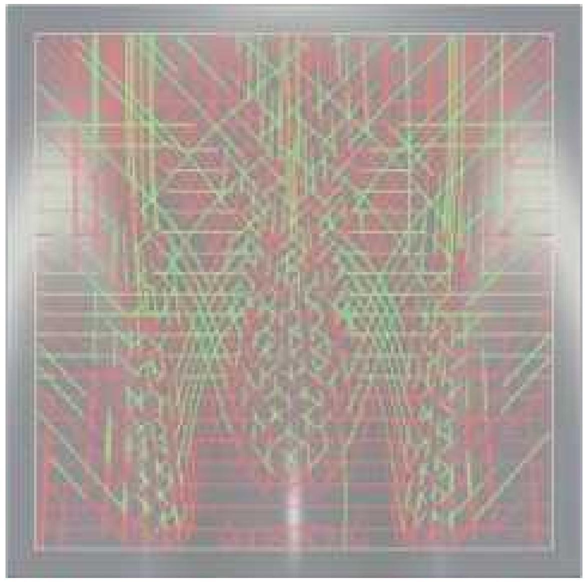

Figure 4.14: LLO baseline defect distribution for the complete Quartic laminate.

which span four different boundary types, exhibit less severe stacking because the different spatial extents of their boundaries naturally prevent complete overlap.

Because the two severity kernels compute defect interaction differently, the same baseline laminate produces two distinct L-IDS values, one under each kernel. Table 4.6 reports both baseline scores. The gravity-based kernel produces a lower baseline score than the stagger shift kernel for the same physical configuration, consistent with the expected behavior that the gravity-based kernel incorporates additional penalty contributions from spatial proximity interactions that the coincident stacking kernel does not capture. Both scores are computed on the same scale, bounded above by 100, but are not directly numerically comparable as absolute values, as discussed in Section 3.5.2.5.

Table 4.6: Baseline L-IDS values for the Quartic laminate under both severity kernels.

Severity Kernel	**Baseline L-IDS**
Stagger shift	90.05
Gravity-based	82.13

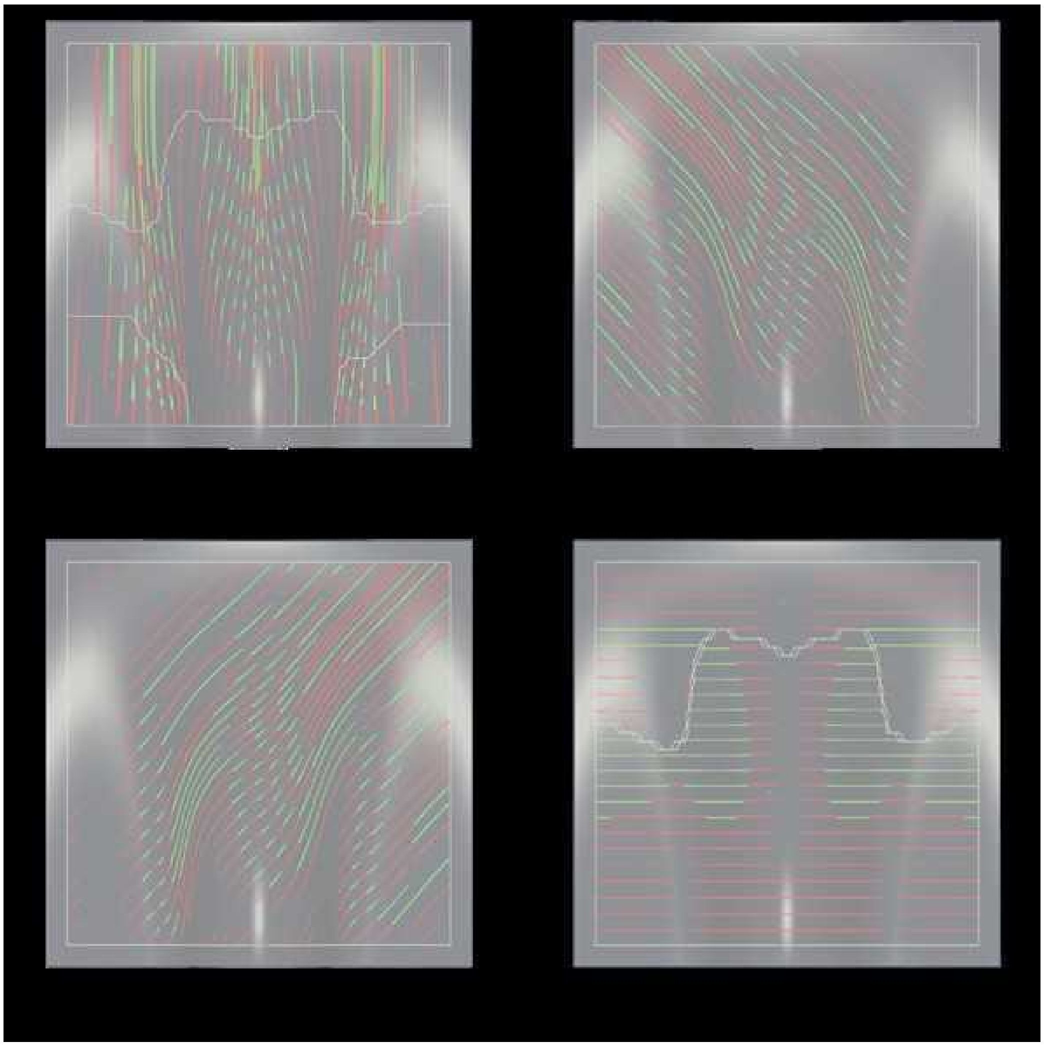

Figure 4.15: Gap and overlap defects disaggregated by fiber orientation for the baseline Quartic laminate. Identical starting points within each angle group produce direct through-thickness defect stacking.

4.5.2 CANDIDATE POOL GENERATION

The candidate pool for LLO was generated using Mode A (Section 3.5.1.2), producing seven candidate scenarios per ply through perpendicular shifts spanning one course width centered on each baseline starting point, with the seven candidates evenly distributed across this range including the baseline scenario at the center. Seven candidates per ply with 11 plies yields a search space of $7^{11} \approx 1.98 \times 10^9$ possible laminate configurations, making exhaustive enumeration infeasible and motivating the cyclic coordinate optimization strategy.

The candidate pool generation required $N_{\text{VCP}} = 66$ VCP evaluations (6 new shifts per ply, with the baseline already computed). Using the subprocess integration, this phase completed in approximately 15–35 minutes depending on ply size and mesh density.[3] Because both severity kernels operate on the same pre-computed defect data, this pool generation was performed once and reused for both optimization runs.

For the gravity-based kernel, an additional neighbor precomputation step was executed prior to optimization. The mesh graph was constructed from the nodal connectivity of the 3,645-element Quartic mesh, and bounded Dijkstra shortest-path searches were performed from each element to populate the neighbor list within the cutoff radius defined by Equation (3.55). The Quartic configuration uses eight 0.5-in. tows per course (Section 4.1), giving a course-width bound of 4 in. The mean element area is $\bar{A}_e \approx 3.03$ sq in. (mean side length $\sqrt{\bar{A}_e} \approx 1.74$ in.), so the mesh-based reach $k \cdot \sqrt{\bar{A}_e} \approx 6.96$ in. exceeds the course-width bound, and the cutoff radius is governed by the mesh resolution. The precomputation produced an average of approximately 32 neighbors per element (median 31, maximum 43) and completed in under one tenth of a second. This one-time cost is negligible relative to the VCP candidate generation phase and does not significantly affect the total runtime of the gravity-based optimization.

[3]Computation times were obtained on a laptop equipped with an Intel® Core™ Ultra 7 155H processor and 32 GB RAM.

4.5.3 STAGGER SHIFT KERNEL RESULTS

The cyclic coordinate optimization algorithm described in Section 3.5.3 was applied to the Quartic laminate using the stagger shift severity kernel (Section 3.5.2.2). The algorithm converged in three passes—consistent with the theoretical property that monotone improvement on a finite candidate set terminates within a bounded number of iterations—yielding the optimized ply scenario assignment [2/2/6/3/1/1/6/1/5/6/4], where each integer index identifies the selected scenario from the respective ply's candidate pool (index 0 corresponds to the baseline starting point for that ply).

Table 4.7 presents the full stacking sequence with the optimized starting points produced by the stagger shift kernel. Table 4.8 summarizes the L-IDS improvement achieved.

Table 4.7: Stacking sequence and optimized starting points for the Quartic model under the stagger shift kernel.

Ply	**Fiber Angle**	**Ply Boundary**	**Optimized Starting Point (in.)**
1	90°	Partial	(−3.586, −8.005, 32.841)
2	45°	Full	(4.130, 0.032, 28.897)
3	−45°	Full	(4.811, 1.328, 27.993)
4	0°	Full	(6.037, 0.397, 27.810)
5	0°	Bottom left corner	(36.173, −35.101, −3.064)
6	0°	Top left corner	(−34.489, −33.115, −3.419)
7	0°	Partial complement	(−1.149, 35.074, 13.339)
8	90°	Partial complement	(−0.593, 35.021, 13.367)
9	−45°	Full	(5.507, −0.017, 28.296)
10	45°	Full	(4.764, 0.596, 28.364)
11	90°	Partial	(−4.288, −9.948, 33.384)

Table 4.8: Stagger shift kernel L-IDS comparison for the Quartic model.

Metric	**Baseline**	**Optimized**
L-IDS	90.05	92.59
L-IDS improvement	+2.54	
Scenario designator	[0/0/0/0/0/0/0/0/0/0/0]	[2/2/6/3/1/1/6/1/5/6/4]
Passes to convergence	N/A	3

The optimized defect distribution under the stagger shift kernel is shown in Figure 4.16. Comparison with the baseline in Figure 4.14 reveals a more uniform spatial distribution of gaps and overlaps, with the high-density stacking concentrations characteristic of the baseline visibly reduced.

Figure 4.16: LLO optimized defect distribution for the Quartic laminate under the stagger shift severity kernel. Defect concentrations are redistributed relative to the baseline, reducing through-thickness stacking.

The staggering effect is most pronounced when examined by orientation, as shown in Figure 4.17. For the +45° ply pair (Plies 2 and 10), the optimized starting points are spatially offset from one another and from the original common baseline starting point, shifting the course patterns relative to one another so that the gap and overlap locations in Ply 2 no longer coincide with those of Ply 10. An analogous staggering is achieved for the −45° pair (Plies 3 and 9).

4.5.4 GRAVITY-BASED KERNEL RESULTS

The same cyclic coordinate optimization algorithm was then applied to the Quartic laminate using the gravity-based severity kernel (Section 3.5.2.3), drawing from the same

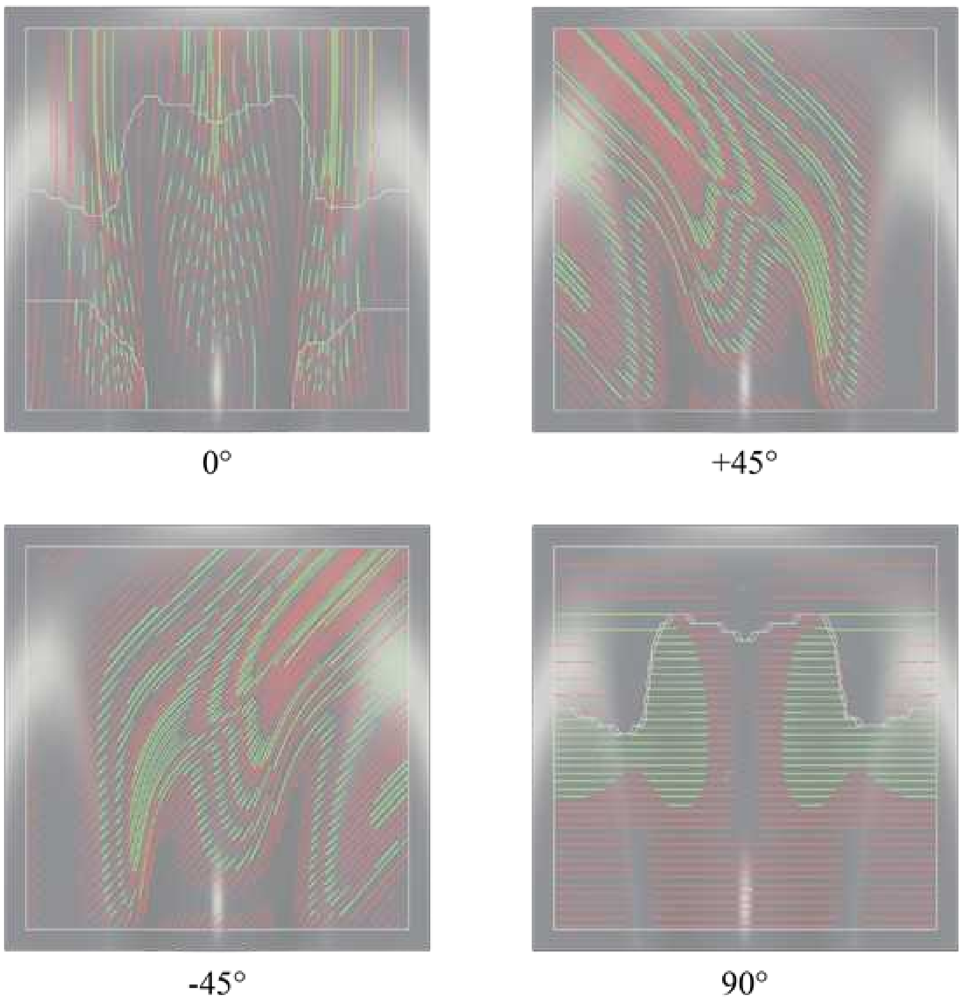

Figure 4.17: Gap and overlap defects disaggregated by fiber orientation for the stagger-shift-optimized Quartic laminate. Starting-point staggering redistributes defect locations between plies of identical orientation, reducing through-thickness coincidence.

candidate pool used in Section 4.5.3. The only differences between the two runs are the severity kernel used to compute the element-level severity $\phi_{m,n}$ in Equation (3.61) and the one-time neighbor precomputation step described in Section 4.5.2.

The algorithm converged in four passes, yielding the optimized ply scenario assignment [3/4/6/5/1/1/4/2/3/5/5]. Table 4.9 presents the full stacking sequence with the optimized starting points produced by the gravity-based kernel. Table 4.10 summarizes the L-IDS improvement achieved.

Table 4.9: Stacking sequence and optimized starting points for the Quartic model under the gravity-based kernel.

Ply	**Fiber Angle**	**Ply Boundary**	**Optimized Starting Point (in.)**
1	90°	Partial	(−3.236, −7.421, 32.564)
2	45°	Full	(3.482, −0.617, 29.401)
3	−45°	Full	(4.811, 1.328, 27.993)
4	0°	Full	(6.892, 1.183, 27.518)
5	0°	Bottom left corner	(36.173, −35.101, −3.064)
6	0°	Top left corner	(−34.489, −33.115, −3.419)
7	0°	Partial complement	(−1.812, 34.396, 13.602)
8	90°	Partial complement	(0.084, 35.347, 13.241)
9	−45°	Full	(6.218, 0.704, 27.847)
10	45°	Full	(5.443, 1.275, 28.108)
11	90°	Partial	(−4.954, −10.602, 33.712)

Table 4.10: Gravity-based kernel L-IDS comparison for the Quartic model.

Metric	**Baseline**	**Optimized**
L-IDS	82.13	86.47
L-IDS improvement	+4.34	
Scenario designator	[0/0/0/0/0/0/0/0/0/0/0]	[3/4/6/5/1/1/4/2/3/5/5]
Passes to convergence	N/A	4

The optimized defect distribution under the gravity-based kernel is shown in Figure 4.18. As with the stagger shift result, the high-density stacking concentrations of the baseline are reduced, though the specific spatial redistribution of defects reflects the gravity-based kernel's sensitivity to in-plane proximity in addition to direct through-thickness coincidence.

Figure 4.18: LLO optimized defect distribution for the Quartic laminate under the gravity-based severity kernel. Defect concentrations are redistributed to minimize both direct stacking and near-miss proximity interactions.

The orientation-disaggregated view in Figure 4.19 illustrates how the gravity-based kernel redistributes defects within each fiber angle group. The optimized starting points achieve not only direct offset between same-orientation ply pairs but also lateral dispersion that separates defects by more than one element in-plane where the candidate pool permits, reflecting the gravity-based kernel's penalty on near-miss proximity.

4.5.5 COMPARATIVE EVALUATION

With both optimization results obtained, the two severity kernels can be evaluated directly against one another on the same baseline laminate and candidate pool. This comparison isolates the effect of the severity kernel from all other framework components, since the candidate pool, optimization algorithm, convergence criterion, frequency penalty, and margin weighting are identical across the two runs.

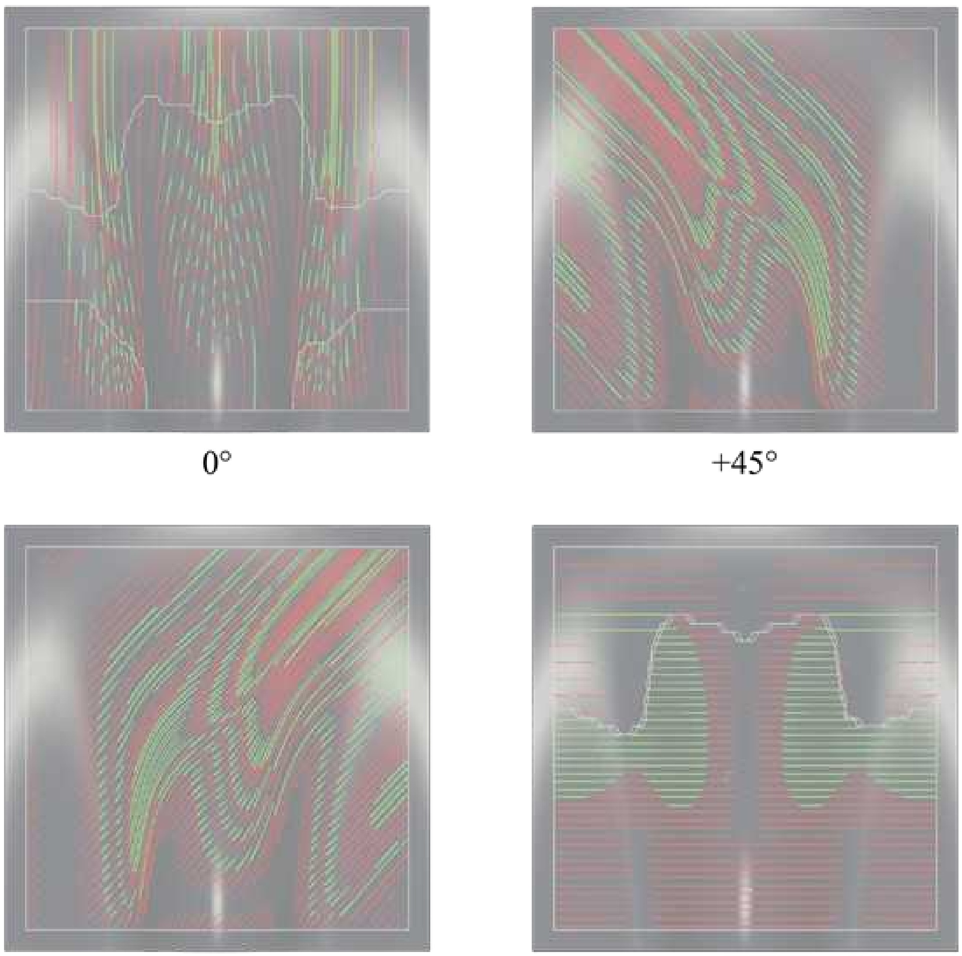

0° +45°

-45° 90°

Figure 4.19: Gap and overlap defects disaggregated by fiber orientation for the gravity-based-optimized Quartic laminate. Starting-point redistribution addresses both direct stacking and in-plane proximity between defects in same-orientation plies.

Table 4.11 presents a side-by-side summary of the optimization outputs under both kernels. The reported metrics include the optimized ply scenario designators, the L-IDS scores under each kernel's own scoring (so that improvements are measured within a consistent framework), the number of passes to convergence, and the iterative optimization runtime. The magnitude of improvement under both kernels is modest in absolute terms; the structural reasons for this are addressed in the discussion of Section 4.5.6, which contextualizes the Quartic model as a constrained staggering scenario.

Table 4.11: Side-by-side comparison of LLO results under the stagger shift and gravity-based severity kernels. Scores are reported under each kernel's own scoring formulation.

Metric	**Stagger Shift Kernel**	**Gravity-Based Kernel**
Baseline L-IDS	90.05	82.13
Optimized L-IDS	92.59	86.47
L-IDS improvement	2.54	4.34
Scenario designator	[2/2/6/3/1/1/6/1/5/6/4]	[3/4/6/5/1/1/4/2/3/5/5]
Passes to convergence	3	4
Iterative optimization time	12.4 s	38.7 s

Beyond the aggregate score comparison, the two kernels can be evaluated on whether they agree or disagree on individual ply assignments. Table 4.12 reports the optimized candidate pool index selected for each ply under both kernels, with a column indicating whether the two kernels selected the same scenario. Agreement on a particular ply indicates that both kernels identify the same candidate as optimal given the rest of the laminate, while disagreement signals that the two kernels are sensitive to different aspects of the defect arrangement at that ply.

A second informative comparison is the evaluation of each kernel's optimized laminate under the *other* kernel's scoring function. This cross-evaluation reveals how well a laminate optimized under one interaction model performs when judged against the other model's criteria. Table 4.13 reports these cross-scores.

Table 4.12: Per-ply comparison of optimized candidate selections under the two severity kernels.

Ply	Fiber Angle	Stagger Shift Index	Gravity-Based Index	Agreement
1	90°	2	3	No
2	45°	2	4	No
3	−45°	6	6	Yes
4	0°	3	5	No
5	0°	1	1	Yes
6	0°	1	1	Yes
7	0°	6	4	No
8	90°	1	2	No
9	−45°	5	3	No
10	45°	6	5	No
11	90°	4	5	No

Table 4.13: Cross-evaluation of optimized laminates under both severity kernels. Each row reports the L-IDS of the indicated optimized laminate when scored by both kernels.

Optimized Laminate	Stagger Shift L-IDS	Gravity-Based L-IDS
Baseline	90.05	82.13
Stagger shift optimized	92.59	84.21
Gravity-based optimized	92.18	86.47

4.5.6 DISCUSSION

The LLO results on the Quartic benchmark support several conclusions about the unified framework and the comparative behavior of its two severity kernels.

First, the cyclic coordinate optimization algorithm consistently improves the laminate-level score under both kernels: the monotonic convergence guarantee established in Section 3.5.3.2 is confirmed empirically, with measurable L-IDS improvements achieved in a small number of passes in both runs. This confirms the claim in Section 3.5.3.2 that the convergence properties of the algorithm are independent of the severity kernel's internal structure: both the coincident-stacking and spatially-decaying kernels admit the same deterministic optimization procedure with the same finite-convergence guarantee.

Second, the mechanism of improvement is physically interpretable under both kernels as a form of defect staggering, though the two kernels produce subtly different staggering patterns. The stagger shift kernel seeks to eliminate direct element-level coincidence, while the gravity-based kernel additionally seeks to disperse defects across a local neighborhood so that near-miss stacking is also penalized. In regions of the laminate where the candidate pool affords sufficient spatial range, the gravity-based kernel tends to push defects further apart in-plane than the stagger shift kernel, producing a qualitatively different defect distribution even when the number of optimized ply changes is similar.

Third, the agreement and disagreement pattern between the two kernels (Table 4.12) provides insight into which plies are governed by direct stacking versus proximity interactions. Plies where both kernels select the same candidate indicate that direct coincident stacking is the dominant concern for that ply, and eliminating it addresses the proximity concern as well. Plies where the two kernels disagree indicate that the optimal candidate depends on which interaction mode is prioritized, and the choice between kernels becomes a modeling decision tied to the user's structural concerns.

Fourth, the cross-scoring results in Table 4.13 illuminate the relationship between the two kernels' objectives. A laminate optimized under the stagger shift kernel is expected to achieve only a partial improvement under the gravity-based kernel, because it addresses direct stacking but may leave near-miss proximity unreduced. Conversely, a laminate optimized under the gravity-based kernel is expected to achieve most of the available stagger shift improvement while additionally reducing proximity interactions. This asymmetric relationship follows from the strict generalization property established in Equation (3.58): the gravity-based kernel contains the coincident stacking kernel as a special case, so optimization under the more general kernel also addresses the objectives of the more restricted kernel.

Fifth, both algorithms respect the boundary constraints: corner plies and small-boundary plies whose limited spatial extent prevents meaningful shifting are correctly left at or near their baseline positions by both optimizers. This indicates that the candidate pool filtering (boundary containment checks during Mode A construction) combined with the monotonic optimization correctly handles degenerate cases where no meaningful optimization is possible for a particular ply.

The modest magnitude of L-IDS improvement observed under both kernels—2.54 points for the stagger shift kernel and 4.34 points for the gravity-based kernel—is a direct consequence of the structural constraints of the Quartic benchmark rather than a limitation of the LLO framework. The Quartic laminate contains only 11 plies across four fiber orientations, and the ply boundary intersections limit the spatial range over which defects can be redistributed; the combinatorial staggering space available to the optimizer is narrow by construction. The Quartic model presents a relatively constrained staggering scenario because the laminate contains only 11 plies across four fiber orientations, and the ply boundary intersections limit the spatial range over which defects can be redistributed. Laminates with greater ply counts, more orientations, and larger common boundary areas would provide a richer combinatorial space for both staggering and proximity-based redistribution, and are expected to yield proportionally larger LLO improvements under both kernels. The gravity-based kernel, in particular, is expected to produce disproportionately larger gains relative to the stagger shift kernel on larger laminates, because the additional neighborhood interactions it captures scale more rapidly with the size of the design space than the coincident stacking interactions alone. This hypothesis could not be tested on the Quartic model alone and represents a natural direction for future evaluation on larger-scale benchmarks.

Despite these scale limitations, the LLO results on the Quartic model demonstrate both severity kernels as systematic, data-driven replacements for the ad hoc manual staggering currently practiced in the AFP community. The availability of two complementary kernels within a unified framework allows users to select the interaction model that best matches their structural concerns—coincident stacking for applications where direct through-thickness alignment is the primary concern, or gravity-based interaction for applications where spatial concentration of defects across a local neighborhood is also of structural consequence—without requiring changes to the underlying optimization infrastructure, candidate pool, or scoring aggregation.

4.6 INSPECTION DATA VISUALIZATION DEMONSTRATION

The inspection data integration capability described in Section 3.6 provides a unified visualization environment in which as-designed defect predictions from VCP and as-manufactured defect measurements from RIPITx can be compared directly on the tool surface geometry. This section demonstrates this capability using the Quartic model and discusses its role in the iSPP closed-loop workflow.

4.6.1 DEMONSTRATION SETUP

Because the Quartic model is a virtual benchmark and was not physically fabricated, the inspection demonstration employs synthetic data representative of RIPITx output. The synthetic dataset was constructed to replicate the format, spatial density, and defect width statistics typical of profilometer measurements from AFP layup, while intentionally including minor spatial offsets relative to the VCP predictions to exercise the coordinate transformation and visualization comparison features. Verification of RIPITx measurement accuracy is outside the scope of this work; the system has been independently characterized in prior studies [59].

4.6.2 VISUALIZATION RESULTS

Figure 4.20 presents the integrated visualization for Ply 1 of the Quartic model. The display applies a minimum width threshold of 0.050 in. to filter sub-threshold measurement noise and a magnification factor of five to make defect line segments legible at the full-scale part geometry. Inspection data points (yellow for detected gaps, purple for detected overlaps) are overlaid on the VCP-predicted course geometry (course centerlines as dashed lines, predicted gaps as green polygons, predicted overlaps as red polygons).

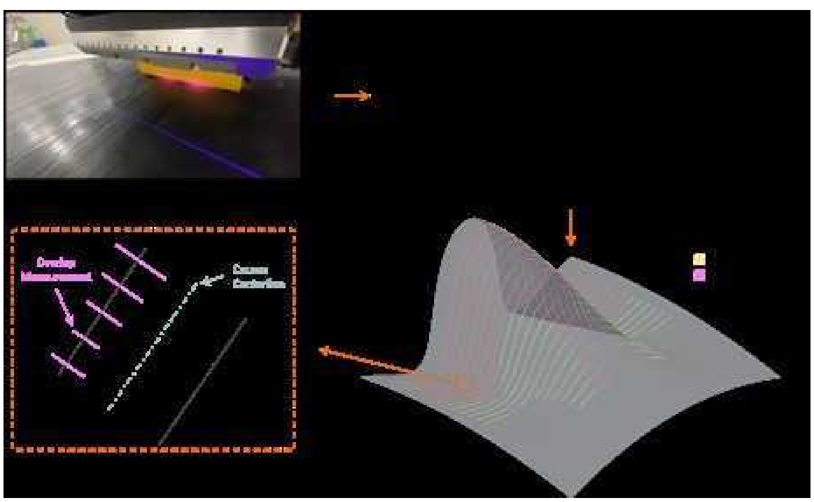

Figure 4.20: Integrated inspection visualization for Quartic model Ply 1. Yellow line segments indicate RIPITx-detected gaps; purple indicates detected overlaps. Green and red polygons show VCP-predicted gaps and overlaps, respectively. Width threshold: 0.050 in.; magnification factor: 5×.

Several features of the visualization illustrate the complementary detection scopes of the two data sources. Inspection data points appear throughout the ply domain, capturing both intra-course defects arising from tow behavior within each course and inter-course defects at course boundaries, while VCP-predicted defects appear as polygons at course boundaries only. The presence of intra-course measurement points with no corresponding VCP predictions, and the presence of VCP inter-course predictions alongside inspection data, confirms the relationship described in Section 3.6.1: the two sources provide overlapping but distinct defect information, with RIPITx additionally capturing intra-course behavior that VCP does not model. The two data sources are not redundant; their inter-course domains overlap, but RIPITx additionally provides as-manufactured measurements of intra-course defects that VCP cannot predict, and VCP provides geometric predictions before manufacturing that RIPITx cannot supply. Together they provide a more comprehensive view of the as-manufactured state than either source alone.

4.6.3 IMPLICATIONS FOR PROCESS REFINEMENT

The simultaneous visualization of predicted and observed defect distributions enables several process improvement analyses that are not possible when inspection data is reviewed independently of the process plan. Clusters of inspection points at consistent locations within a course indicate persistent intra-course defects that may arise from material behavior, machine calibration, or underlying surface contour effects. If these clusters align with regions where VCP predicts elevated inter-course gaps—as would occur, for example, at aggressive curvature transitions characteristic of the Quartic geometry—the combined evidence points to a process parameter adjustment rather than a material or machine issue. Regions where VCP-predicted overlaps from earlier plies coincide with intra-course gap detections from the current ply suggest that built-up surface contours from prior material deposition are disrupting tow behavior, which may indicate that a starting-point adjustment to redistribute the prior-ply overlaps would improve layup quality in subsequent courses.

These analyses require the integrated data environment that CAPP provides. Without a platform that overlays prediction and measurement on the same geometric reference frame, the causal relationships between process parameters, predicted defects, and observed manufacturing behavior remain opaque. The iSPP framework closes this interpretive gap.

4.7 DISCUSSION

The verification and demonstration studies presented in Sections 4.2 through 4.6 collectively demonstrate that each component of the iSPP framework functions as designed and that the components interact to produce improvements not achievable by any single component in isolation. This section synthesizes the results, discusses the characteristics of the IDS scoring behavior that emerged from the case study, and identifies the limitations of the virtual demonstration context.

4.7.1 COMPONENT-LEVEL FINDINGS

The subprocess integration verification established the computational foundation for all subsequent results. The exact agreement between native VCP and CAPP-integrated outputs (Table 4.2) confirms that the subprocess architecture introduces no bias into defect predictions or scores, and the 91.3% reduction in total processing time across 63 ply scenarios transforms the computational feasibility of iterative optimization. Without this efficiency, the PLO survey phases producing 58–69 scenarios per ply and the LLO candidate pool requiring 66 additional VCP evaluations would be impractical within reasonable time budgets.

The HyperX data exchange verification confirmed that structural margins are preserved without distortion through the import process, establishing the data integrity prerequisite for margin-weighted scoring. The presence of negative MoS values in the baseline configuration—86 elements, clustered at the ply-drop region—demonstrates that the Quartic model presents a genuine structural challenge, not merely a synthetic one, and that the DIDS weighting has physical significance for this geometry.

4.7.2 SCORE BEHAVIOR AND DISCRIMINATIVE RANGE

A consistent feature of the scoring results is that all ply scenarios for a given geometry cluster within a relatively narrow portion of the $[0, 100]$ score range. For the Quartic model, the majority of evaluated scenarios score between 88 and 98 on the MIDS scale, with the full spread of achievable improvement across all survey candidates typically spanning 2–8 points for a given ply.

This behavior deserves precise interpretation rather than dismissal as a limitation. The MIDS is a fractional-area metric: a defect's contribution to the penalty is proportional to its spatial footprint relative to the ply area. For the Quartic surface, the surface curvature dictates that gap and overlap defects at course boundaries will occupy a bounded fraction of the ply area across all feasible starting-point configurations. Consequently, the maximum achievable improvement through starting-point optimization is similarly bounded—not by a deficiency in the scoring algorithm, but by a physical property of the geometry and process. A score of 92 versus 98 does not mean the two configurations are nearly equivalent; it means that the second configuration has approximately 6% less ply area covered by threshold-exceeding defects, which represents a real and structurally meaningful improvement in material coverage.

Furthermore, the narrow score range does not impede the optimization algorithms. The PLO algorithm identifies the best achievable configuration within the constrained landscape, and the LLO algorithm correctly selects starting-point combinations that minimize through-thickness stacking even when individual ply score differences are small. Under the stagger shift kernel, the LLO produces a 2.54-point L-IDS improvement; under the gravity-based kernel, which additionally penalizes in-plane defect proximity, the improvement reaches 4.34 points. Both results are entirely attributable to reduction in defect stacking—

a laminate-level property that has no ply-level analog—demonstrating that the scoring architecture captures multi-scale manufacturing quality effects that single-ply metrics would miss entirely. The larger improvement under the gravity-based kernel reflects its sensitivity to near-miss proximity interactions that the stagger shift kernel does not capture, confirming that the choice of severity kernel is a meaningful modeling decision rather than an arbitrary one.

4.7.3 THE ITERATIVE DESIGN-MANUFACTURING LOOP

The DIDS and IDS results in Section 4.3.3 demonstrate the iterative character of the iSPP workflow in its most direct form. The baseline process plan, which does not account for structural criticality, produces a defect distribution that concentrates negative-margin elements at the ply-drop region. The IDS-optimized process plan redistributes defects away from this region, reducing the count of negative-margin elements by 30% without any addition of material. The subsequent HyperX reanalysis identifies residual negative margins that starting-point optimization cannot resolve, and targeted reinforcement plies are added to achieve a fully positive margin configuration.

This sequence mirrors the iterative design process that engineers perform manually today, but with the difference that each step is quantitatively informed by the preceding one. The IDS provides the numerical interface that allows process planning decisions to be evaluated in terms of structural consequences, and the HyperX–CAPP digital thread provides the data pathway that makes this evaluation automated rather than judgment-dependent.

4.7.4 LIMITATIONS OF THE VIRTUAL DEMONSTRATION

The Quartic model demonstration is entirely computational, which appropriately delimits the claims that can be drawn from the results. No physical part was fabricated, and the inspection demonstration employs synthetic RIPITx-format data to exercise the data import, coordinate transformation, and visualization capabilities of CAPP rather than to characterize measurement accuracy. Accordingly, the demonstration confirms the functional correctness and internal consistency of the iSPP framework components — that each module operates as formulated and that the integrated workflow produces physically interpretable outputs — but does not address the fidelity of VCP defect predictions relative to an actual manufactured part. Physical validation of the framework against a fabricated test article, including correlation of VCP predictions with real RIPITx measurements, represents a natural and well-defined direction for future work.

Within the virtual demonstration context, the PLO and LLO results are subject to the accuracy of VCP's defect prediction model. The subprocess verification confirmed that CAPP faithfully reproduces VCP's outputs, but it does not address whether those outputs accurately represent the defects that would appear on a physical part. Industry experience and prior validation studies [202] support the use of VCP as a reliable predictor of gap and overlap geometry for standard AFP processes, but surface-specific validation remains good practice before relying on optimization results for production planning.

4.7.5 SUMMARY

Table 4.14 summarizes the quantitative outcomes of the Quartic model case study across all iSPP framework components.

The case study demonstrates that the iSPP framework components are individually correct, mutually consistent, and collectively capable of improvements that span from ply-level defect redistribution through laminate-level stacking optimization to structural margin improvement. Chapter 5 summarizes the contributions of this work and identifies directions for future development.

Table 4.14: Summary of Quartic model verification and demonstration results across iSPP framework components.

Component	Key Result	Section
VCP subprocess integration	Exact output agreement; 91.3% time reduction	4.2
HyperX data exchange	Zero-loss import of 3,645 elements, MoS data	4.3.2
MIDS scoring	2.73-point discrimination between Ply 1 scenarios	4.3.1
DIDS/IDS optimization	≈30% reduction in negative-margin elements	4.3.3
Structural reinforcement	All margins positive; min MoS = 0.437	4.3.3
PLO (Ply 1, Ply 2)	IDS: 98.51→98.52; 97.94→98.05	4.4
LLO stagger shift	L-IDS: 90.05→92.59 (+2.54 points)	4.5
LLO gravity-based	L-IDS: 82.13→86.47 (+4.34 points)	4.5
Inspection visualization	Complementary VCP/RIPITx scope demonstrated	4.6

CHAPTER 5

CONCLUSIONS AND FUTURE WORK

The case study presented in Chapter 4 verified and demonstrated each component of the iSPP framework through systematic application to the Quartic benchmark surface. Taken together, the results confirm that the framework's components are individually correct, mutually consistent, and collectively capable of improvements that span from ply-level defect redistribution through laminate-level stacking optimization to structural margin improvement. This sequence of quantitatively connected decisions is currently executed in isolation, if at all, in current practice. The present chapter closes the dissertation by synthesizing these contributions, situating them within the broader research landscape, and identifying the directions in which this work should be extended.

Section 5.1 summarizes the seven contributions of this work, tracing each from its motivating gap through its technical realization and demonstrated outcome. Section 5.2 situates these contributions within the scholarly record, maps them explicitly to the four research gaps identified in Chapter 2, and positions them relative to the Smart AFP vision and prior work from the neXt research group and the broader AFP community. Section 5.3 identifies directions for continued development, organized into near-term extensions that follow directly from the present results and longer-term directions that require physical manufacturing data, expanded geometries, or new enabling capabilities.

5.1 SUMMARY OF CONTRIBUTIONS

AFP process planning has historically been performed as a series of largely independent activities: structural designers optimize laminates for idealized fiber paths, process planners select starting points and layup strategies through experience-driven trial and error, and inspection data is reviewed in isolation after manufacturing is complete. The consequence of this fragmentation is that manufacturing decisions carry structural implications that are never quantified at the time they are made, and that defect information accumulated during production does not systematically inform future planning cycles. This dissertation addressed these limitations through the development of the iSPP framework, an integrated platform that connects design, process planning, manufacturing, and inspection through bidirectional data exchange and provides the scoring, optimization, and visualization tools necessary to make process planning decisions that are simultaneously manufacturable and structurally informed.

CONTRIBUTION 1: COMPUTER-AIDED PROCESS PLANNING SOFTWARE PLATFORM

The first and foundational contribution is the CAPP software, a purpose-built process planning environment that serves as the integration hub of the iSPP framework. CAPP orchestrates bidirectional data exchange with HyperX for structural margin data, VCP for defect prediction, path generation, and NC code output, and RIPITx for as-manufactured inspection measurements. Beyond data exchange, CAPP provides the computational environment within which the scoring algorithms and optimization routines of subsequent contributions operate, and a visualization interface that allows process planners to inspect defect distributions, structural margins, and inspection overlays on the tool surface geometry. The platform is the enabling architecture without which the remaining contributions would have no integrated operational context.

CONTRIBUTION 2: VCP SUBPROCESS INTEGRATION

The second contribution is the CAPP–VCP subprocess integration architecture, which provides the computational infrastructure on which the optimization contributions depend. By maintaining a persistent VCP process in memory and modifying only the specific ply under evaluation at each iteration, the subprocess architecture eliminates the complete laminate regeneration overhead that serial socket-based integration required in prior work. The practical consequence is a 91.3% reduction in processing time for a representative set of 63 ply scenarios (from 4.54 hours to 23.65 minutes) transforming the computational budget available for optimization. Without this efficiency, the PLO survey phases and the LLO candidate pool generation, which together require hundreds of individual VCP evaluations per study, would be impractical within the time constraints of an engineering workflow.

CONTRIBUTION 3: DEFECT SCORING METHODOLOGY

The third contribution is a unified defect scoring methodology comprising three novel metrics that operate at progressively broader scopes. The MIDS quantifies ply-level manufacturing quality by aggregating frequency and severity penalties across all defect types, each weighted by user-specified thresholds and importance factors calibrated through an AHP pairwise comparison process, into a single score on a bounded $[0, 100]$ scale. The DIDS extends this framework by incorporating element-level structural margin data from HyperX, weighting defect penalties by the structural criticality of the regions in which they occur so that defects in low-margin zones contribute disproportionately to the score. The IDS combines MIDS and DIDS through user-defined weights, enabling explicit trade-off analysis between manufacturing efficiency and structural performance. Together these three metrics provide the quantitative interface that allows process planning decisions to be evaluated in terms of their structural consequences, a capability that did not previously exist in AFP process planning practice.

CONTRIBUTION 4: PLY-LEVEL OPTIMIZATION

The fourth contribution is the PLO algorithm, a two-phase method for identifying the starting point and layup strategy configuration that maximizes the IDS for a given ply. In the first phase, a global survey evaluates the IDS at a structured grid of candidate starting points spanning the ply domain, providing coarse coverage sufficient to identify high-scoring regions and avoid convergence to local optima. The highest-scoring survey candidates are retained as seeds for the second phase, in which a perpendicular-shift refinement locally explores the neighborhood of each seed along the course-width direction. The algorithm makes no smoothness assumptions about the IDS landscape, a necessary design choice, as the Quartic case study confirmed that the score profile is discontinuous as a function of perpendicular shift distance and therefore incompatible with gradient- or interpolation-based refinement methods. Applied to the Quartic benchmark, PLO identified improved starting configurations for both demonstration plies, with the IDS increasing from 97.94 to 98.05 for the more geometrically complex ply.

CONTRIBUTION 5: LAMINATE-LEVEL OPTIMIZATION

The fifth contribution is the LLO algorithm, which extends process optimization through the laminate thickness by selecting combinations of ply scenarios whose through-thickness defect distributions are spatially offset, minimizing the accumulation of stacked defects that produce the most structurally damaging laminate conditions. LLO employs a cyclic coordinate optimization strategy that searches the combinatorial space of ply scenario assignments through sequential per-ply optimization passes that converge in three to four iterations. Two complementary severity kernels characterize defect interaction differently: the stagger shift kernel penalizes direct spatial coincidence of same-type defects across adjacent plies, while the gravity-based kernel, formulated by analogy to Newton's law of universal gravitation, additionally penalizes near-miss proximity interactions within a geodesic neighborhood weighted by defect density. Applied to the Quartic laminate, LLO produced L-IDS improvements of 2.54 points under the stagger shift kernel and 4.34 points under the gravity-based kernel, with the larger improvement under the gravity-based kernel reflecting its sensitivity to proximity interactions that the stagger shift kernel does not capture.

CONTRIBUTION 6: INSPECTION DATA INTEGRATION

The sixth contribution is the integration of as-manufactured inspection data from RIPITx into the CAPP environment, closing the feedback loop between production outcomes and process planning. CAPP imports RIPITx measurement files, applies the coordinate transformation required to register inspection data to the tool surface geometry, and overlays detected defects (both intra-course and inter-course) against the VCP-predicted course geometry and as-designed defect predictions on a shared geometric reference frame. This co-visualization enables process planners to identify systematic discrepancies between predicted and observed defect distributions, to distinguish defects attributable to process parameters from those arising from material behavior or machine calibration, and to carry observations from one manufacturing cycle forward into the next planning iteration. The capability transforms inspection data from an isolated quality-control record into an active component of the process planning workflow.

CONTRIBUTION 7: QUARTIC BENCHMARK VERIFICATION AND DEMONSTRATION

The seventh contribution is the systematic verification and demonstration of the complete iSPP framework through application to the Quartic benchmark surface, a doubly-curved saddle geometry developed by Collier Aerospace as a representative and geometrically challenging AFP process planning test case. The case study exercises each framework component in sequence—subprocess integration, defect scoring, structural data exchange, ply-level optimization, laminate-level optimization, and inspection visualization—under a consistent set of process planning parameters and a fixed laminate configuration, confirming that the components operate as formulated individually and produce physically interpretable, mutually consistent outputs when integrated. The Quartic demonstration establishes that the iSPP framework is not merely a collection of independently functional modules but a coherent system capable of the full lifecycle-connected decision sequence that motivated its development.

BROADER SIGNIFICANCE

Collectively, the seven contributions advance AFP process planning from a siloed, experience-dependent activity toward a smart, lifecycle-integrated discipline consistent with the Smart AFP vision articulated by the neXt research group. The modular architecture of the iSPP framework ensures that improvements to individual components, driven by advances in commercial simulation software, new inspection sensor technologies, or more capable optimization methods, can be incorporated without fundamental redesign of the integration infrastructure. By semi-automating decisions that currently require experienced engineers and embedding structural performance awareness directly into the manufacturing objective, the framework lowers the expertise threshold for AFP process planning and positions the technology for broader industrial application beyond aerospace, including automotive, maritime, and wind energy sectors where the cost barriers to AFP adoption remain most significant.

5.2 SITUATION OF WORK

5.2.1 ADDRESSING THE RESEARCH GAPS

Chapter 2 identified four gaps that collectively characterize the state of AFP process planning prior to this work: the open-loop, siloed relationship between the design, process planning, manufacturing, and inspection pillars of the AFP lifecycle (Gap 1); the decoupling of structural and manufacturing metrics in process planning decisions (Gap 2); the absence of efficient, integrated optimization workflows capable of operating within practical engineering time budgets (Gap 3); and the lack of a closed-loop mechanism for incorporating as-manufactured inspection data into the process planning cycle (Gap 4). Figure 5.1 maps the seven contributions of this dissertation to these four gaps.

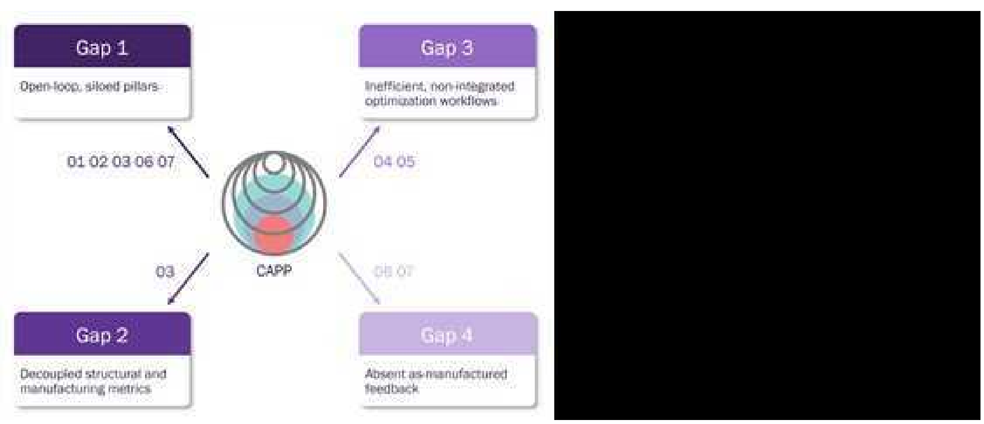

Figure 5.1: Mapping of the seven iSPP framework contributions to the four research gaps identified in Chapter 2.

5.2.1.1 GAP 1: OPEN-LOOP, SILOED LIFECYCLE PILLARS

Gap 1 reflects the historical absence of bidirectional data exchange between the design, process planning, manufacturing, and inspection stages of the AFP lifecycle. Prior to this work, each stage operated with its own tools, data formats, and decision criteria, with no systematic mechanism for propagating information from one stage to another.

Contributions 1 through 3 and Contributions 6 and 7 collectively address this gap. The CAPP platform (Contribution 1) establishes a central integration hub through which design intent, manufacturing predictions, and inspection measurements are connected under a unified data model. The VCP subprocess integration (Contribution 2) provides the computational backbone that makes iterative, simulation-driven process planning feasible within this hub. The defect scoring methodology (Contribution 3) provides the shared quantitative language through which manufacturing outcomes can be evaluated against design requirements. The RIPITx inspection integration (Contribution 6) closes the loop by returning as-manufactured data to the planning environment. The Quartic benchmark demonstration (Contribution 7) confirms that this integration operates without data loss or geometric distortion across all connected tools.

The iSPP framework is grounded in the broader neXt research group's vision of Smart AFP [11], which frames the long-term objective as a lifecycle-integrated, data-driven manufacturing system capable of autonomous adaptation. This dissertation represents a concrete realization of that vision at the process planning level, delivering the semi-automated integration infrastructure that prior reviews of the field identified as a persistent gap [7, 19, 51].

Limitation: The iSPP framework was demonstrated on a virtual test case using VCP-predicted defect data. Physical correlation of VCP-predicted defects against a fabricated part was not performed, and the results therefore represent verification of functional correctness and internal consistency rather than experimental validation. Physical validation against a fabricated test article, discussed in Section 5.3, is the natural next step.

5.2.1.2 GAP 2: DECOUPLED STRUCTURAL AND MANUFACTURING METRICS

Gap 2 reflects the absence of a process planning objective that accounts for the structural consequences of manufacturing decisions. Prior to this work, starting point and layup strategy selection was performed with respect to manufacturing quality alone—minimizing defect area or count—without reference to whether the defects that remained were located in structurally critical or non-critical regions of the laminate. Contribution 3 addresses this gap directly through the DIDS and IDS formulations, which incorporate element-level MoS data from HyperX as a spatial weighting on the defect penalty. The most direct predecessor to this capability in the group's prior work is the design-for-manufacturing framework developed in [198, 200], which established the conceptual linkage between structural margin data and manufacturing analysis in the AFP context. The DIDS/IDS formulation in this dissertation operationalizes that linkage by embedding it directly into a differentiable scoring objective that drives optimization. The structural feedback loop demonstrated in Section 4.3, in which IDS-driven starting-point optimization reduced the count of negative-margin elements by approximately 30% and enabled targeted local reinforcement that brought all margins positive, confirms that this coupling produces physically meaningful structural consequences, not merely numerical ones [162].

5.2.1.3 GAP 3: INEFFICIENT, NON-INTEGRATED OPTIMIZATION WORKFLOWS

Gap 3 reflects the absence of optimization methods capable of improving process planning decisions within the time constraints of an industrial workflow. The 91.3% reduction in per-evaluation processing time delivered by the VCP subprocess integration (Contribution 2) is the enabling precondition for Contributions 4 and 5: without it, the survey phases of PLO and the candidate pool generation of LLO, which together require hundreds of VCP evaluations, would require multi-day computation and would be impractical in any iterative design context. With it, both optimization studies were completed within a single working session on standard laptop hardware.[1]

[1]Computation times were obtained on a laptop equipped with an Intel® Core™ Ultra 7 155H processor and 32 GB RAM.

The PLO algorithm (Contribution 4) advances beyond the Bayesian optimization approach for starting-point selection explored by Brasington et al. [178], which assumed a smooth objective landscape. The Quartic case study demonstrated that the IDS landscape is in fact discontinuous as a function of perpendicular shift distance, a finding that invalidates interpolation-based methods and motivates the sampling-based two-phase strategy implemented here. The LLO algorithm (Contribution 5) addresses a scope of optimization not previously treated in the group's published work: through-thickness stagger optimization at the laminate level. The combinatorial through-thickness stacking problem was identified as a critical open direction in the prior theses of Swingle [76] and Patel [77], both of which developed foundational elements of the LLO framework that this dissertation synthesizes and extends into the full iSPP context. The earlier CAPP process planning automation work of Halbritter [53, 203, 213] and the integrated process planning metrics framework of Brasington [22, 200] established the process planning infrastructure and metric foundations on which the PLO and LLO implementations rest.

Limitation: The PLO and LLO algorithms were demonstrated on the Quartic laminate, which contains 11 plies across four fiber orientations. The constrained stagger space of this geometry limits the magnitude of LLO improvement achievable through starting-point variation alone; laminates with greater ply counts and larger common ply boundary areas are expected to yield proportionally larger improvements as the combinatorial space of distinct stagger configurations grows with ply count. The framework architecture is geometry-independent and has been designed to scale to larger laminates without structural modification. Additionally, the LLO severity computation considers same-type defect stacking independently; cross-type interactions between, for example, gap stacking and overlap stacking are excluded because no established quantitative relationships exist in the literature to inform their relative weighting. The scoring architecture is extensible to cross-type interaction terms as the relevant experimental data becomes available, and pursuing this extension is identified as a near-term research direction in Section 5.3.

5.2.1.4 GAP 4: ABSENT AS-MANUFACTURED FEEDBACK

Gap 4 reflects the treatment of inspection as a post-production quality record rather than an active input to the process planning cycle. Prior work in the group on AFP inspection focused on automated defect detection and characterization from profilometry data [60, 71, 195] and on quantifying the structural consequences of as-manufactured defects [162]. Contribution 6 of this dissertation takes the next logical step: integrating the inspection data stream directly into CAPP so that as-manufactured defect distributions can be co-visualized against VCP predictions on a shared geometric reference frame. This co-visualization enables process planners to identify systematic discrepancies between predicted and observed defect patterns and to carry those observations forward into subsequent planning iterations, transforming inspection from an isolated audit into an active component of a continuous improvement cycle.

Limitation: The RIPITx integration was exercised using synthetic data formatted to match the inspection system's output specification; the fidelity and detection accuracy of the RIPITx system itself were not characterized in this work, which focused on the integration architecture rather than sensor performance. Characterization of RIPITx detection fidelity in the context of the iSPP framework is contingent on the physical validation study identified in Section 5.3.

5.2.2 POSITIONING RELATIVE TO PRIOR WORK

The iSPP framework occupies a distinct position relative to the broader AFP process planning literature reviewed in Chapter 2. Existing optimization approaches in the field address individual aspects of process planning—path planning geometry, defect prediction fidelity, or structural analysis—without integrating these aspects into a unified decision framework [7, 19]. The iSPP framework is, to the author's knowledge, the first process planning system to simultaneously connect structural margin data, manufacturing simulation, ply-level optimization, laminate-level stacking optimization, and as-manufactured inspection feedback through a single software platform, operating with commercially available industrial tools in an architecture that requires no proprietary solver or specialized hardware. The modular design ensures that the framework can accommodate advances in any of its constituent tools without requiring fundamental redesign of the integration logic.

Within the neXt research group, this dissertation synthesizes and extends a body of work spanning process planning automation [52, 53, 203], defect scoring and optimization [22, 178, 200], through-thickness stacking analysis [76, 77, 183], and inspection data integration [71, 195].

5.3 FUTURE WORK

The iSPP framework establishes a foundation on which a substantial program of future work can be built. The directions identified here are organized into near-term extensions that follow directly from the present results and require no fundamental change to the framework architecture, and longer-term directions that require physical manufacturing data, expanded computational resources, or new enabling capabilities before they can be fully realized.

5.3.1 NEAR-TERM EXTENSIONS

5.3.1.1 PHYSICAL VALIDATION AGAINST A FABRICATED TEST ARTICLE

The most immediate and consequential extension of this work is physical validation of the iSPP framework against a fabricated test article. The Quartic benchmark case study demonstrated that the framework's components are individually correct and collectively consistent within the virtual simulation environment, but all results represent comparisons of VCP predictions against other VCP predictions rather than against measured physical outcomes. Physical validation would involve fabricating a test laminate on an AFP machine, importing the resulting RIPITx inspection data into CAPP using real rather than synthetic measurements, and comparing the as-manufactured defect distributions against the VCP predictions for the same process planning scenario. This comparison would quantify the predictive fidelity of VCP, a prerequisite for making process planning decisions whose structural implications are grounded in physical reality rather than simulation alone. Physical validation would also exercise the full closed-loop capability of the framework in a way that the virtual demonstration could not, enabling assessment of whether the optimization improvements identified by PLO and LLO survive the transition from predicted to as-manufactured defect distributions.

5.3.1.2 TOW COUNT AS AN OPTIMIZATION VARIABLE

The current PLO formulation treats tow count (the number of tows in a single course) as a fixed process planning parameter specified by the user prior to optimization. In practice, tow count directly influences course width, which in turn determines the spatial frequency of gap and overlap defects at course boundaries. Incorporating tow count as a discrete optimization variable within PLO would expand the feasible process planning space and enable the algorithm to identify configurations that reduce defect occurrence through geometric means rather than just through starting-point and layup strategy selection. The subprocess architecture already supports per-ply tow count specification through the VCP interface, so the primary development required is an extension of the PLO design vector and objective function to treat tow count as a co-optimized variable alongside starting point and layup strategy.

5.3.1.3 CROSS-TYPE DEFECT INTERACTIONS IN LLO

The LLO severity kernels developed in this dissertation compute same-type defect stacking independently: gap accumulation and overlap accumulation are penalized separately, and no interaction term captures the combined effect of stacked gaps and stacked overlaps in the same through-thickness column of elements. This simplification was adopted because no established quantitative relationships exist in the literature to inform the relative structural weighting of cross-type interactions. As experimental data on the combined effects of gap and overlap stacking becomes available the LLO scoring architecture can be extended to include cross-type interaction kernels without structural modification to the cyclic coordinate optimization framework. Establishing these interaction weightings experimentally is itself a research direction of significant practical value, as it would provide the empirical foundation that the cross-type extension requires.

5.3.1.4 AS-MANUFACTURED SCORING

The MIDS, DIDS, and IDS formulations developed in this dissertation operate on VCP-predicted defect data, producing scores that characterize the expected manufacturing quality of a process planning scenario prior to fabrication. A natural extension is to apply the same scoring framework directly to as-manufactured defect data imported from RIPITx, producing measured variants of each metric—denoted m-MIDS, m-DIDS, and m-IDS respectively—that characterize the actual quality of a fabricated ply or laminate rather than its predicted quality. The m-MIDS would quantify the as-manufactured frequency and severity of defects detected by RIPITx; the m-DIDS would weight those measured defects by the structural margins of the regions in which they occur; and the m-IDS would combine both into a single as-manufactured quality score directly comparable to its predicted counterpart. The difference between the predicted and measured scores for the same ply scenario would provide a quantitative measure of prediction fidelity, enabling systematic identification of process planning scenarios for which VCP predictions are reliable and those for which they require calibration. This extension requires no new algorithmic development beyond connecting the existing RIPITx import pipeline to the existing scoring module, making it among the most accessible near-term extensions available.

5.3.2 LONGER-TERM DIRECTIONS

5.3.2.1 EXPANDED BENCHMARK GEOMETRIES AND LARGER LAMINATES

The Quartic benchmark surface is a well-characterized, geometrically challenging test case, but it represents a single doubly-curved geometry with an 11-ply laminate. Demonstrating the iSPP framework on geometries with different curvature profiles, boundary conditions, and ply counts would establish the generality of the contributions and quantify how optimization improvement scales with laminate complexity. In particular, laminates with greater ply counts and larger common ply boundary areas are expected to yield substantially larger LLO improvements, because the combinatorial space of distinct stagger configurations grows rapidly with ply count and the spatial range over which defect redistribution can occur expands with the laminate area. Applying the framework to a production-representative geometry—such as a fuselage panel or wing skin section—would also provide the scale validation necessary to support industrial adoption.

5.3.2.2 HYBRID OPTIMIZATION: CYCLIC COORDINATE AND METAHEURISTIC METHODS

The cyclic coordinate strategy employed by LLO optimizes each ply's scenario assignment sequentially while holding all others fixed, converging in three to four passes for the Quartic laminate. This approach is computationally efficient and practically effective for the small-to-medium laminates considered here, but it is susceptible to local optima in larger laminates where the interaction structure between plies is more complex: a locally optimal assignment for one ply may preclude a globally better configuration that would require simultaneous adjustment of multiple plies to reach. Incorporating a metaheuristic component, such as simulated annealing or a genetic algorithm applied to the full laminate assignment vector, as a complementary global search phase would allow the optimizer to escape local optima that the sequential coordinate strategy cannot exit. A practical hybrid architecture would use the cyclic coordinate strategy as a fast local refinement step seeded by metaheuristic candidates, combining the convergence efficiency of the former with the global coverage of the latter.

5.3.2.3 MACHINE LEARNING FOR DEFECT PREDICTION ACCELERATION

Each PLO and LLO evaluation currently requires a VCP simulation call to generate defect predictions for a candidate process planning scenario. While the subprocess integration reduced per-evaluation time by 91.3%, VCP simulation remains the dominant computational cost of each optimization study. A surrogate model trained on a library of VCP simulation results could replace direct simulation calls during optimization, reducing per-evaluation cost by orders of magnitude and enabling substantially denser sampling of the process planning space. Machine learning approaches applied to AFP inspection data have demonstrated strong predictive capability in related contexts [71], and extending this capability to defect prediction from process planning parameters is a natural progression. The primary challenge is assembling a training dataset of sufficient coverage and diversity, which could be generated systematically using the subprocess integration as a high-throughput simulation engine prior to surrogate training.

5.3.2.4 REAL-TIME ADAPTIVE PROCESS CONTROL

A longer-horizon extension of the inspection integration is the use of as-manufactured defect data to trigger adaptive modifications to the process plan during an ongoing layup, rather than informing the planning of future parts after the fact. In this architecture, RIPITx measurements from each completed course would be ingested by CAPP in near real time, scored using the m-IDS framework described above, and compared against predicted defect distributions. Systematic deviations above a threshold would trigger automated re-optimization of the remaining ply scenarios using the updated defect landscape. This capability would require tight integration between the CAPP optimization pipeline, the AFP machine controller, and the inspection system at latencies compatible with the machine's layup cycle time; a systems integration challenge that substantially exceeds the scope of the present work but is a long-term target for the Smart AFP vision.

5.3.2.5 BROADER INDUSTRIAL APPLICATION

AFP has historically been concentrated in aerospace applications, where the structural performance requirements and long production runs justify the substantial capital investment and specialized process planning expertise the technology demands. One of the principal barriers to broader industrial adoption is precisely the expertise threshold associated with process planning: selecting starting points, layup strategies, and tow counts that achieve acceptable manufacturing quality requires experienced engineers with deep knowledge of the specific machine, material, and geometry at hand. The iSPP framework directly addresses this barrier by semi-automating the process planning decisions that currently require the domain expertise and by providing the quantitative feedback necessary for less experienced planners to understand and act on the results.

By reducing the upfront intellectual cost of AFP process planning, the iSPP framework

Figure 5.2: The Clemson Composites Center vision for composites application across industrial sectors. The iSPP framework contributes to this vision by lowering the process planning expertise barrier that currently limits AFP adoption outside of aerospace.

lowers the barrier to entry for industries where the structural and weight performance benefits of AFP are compelling but the process planning overhead has historically been prohibitive. As illustrated in Figure 5.2, the Clemson Composites Center envisions AFP technology contributing to sectors including automotive, defense, energy, maritime, and sports equipment manufacturing, in addition to its established aerospace applications. In each of these domains, the ability to generate optimized, structurally informed process plans with reduced reliance on specialized expertise would lower the adoption threshold and expand the range of organizations for which AFP is a viable manufacturing option. Realizing this potential requires not only the process planning tools developed in this dissertation but also expanded validation across the geometries, materials, and production volumes characteristic of each target sector. The iSPP framework's modular, tool-agnostic architecture is well positioned to support such a research program.

5.3.3 CLOSING REMARKS

The work presented in this dissertation began from a straightforward observation: the decisions that determine the structural performance of an AFP-manufactured component are made at the process planning stage, yet process planning has historically operated without the tools to quantify or optimize the structural consequences of those decisions. The iSPP framework developed here addresses that gap by connecting, for the first time in a unified software environment, the structural margin data that characterizes design intent, the manufacturing simulation that predicts what a given process plan will produce, the optimization algorithms that identify the best achievable plan, and the inspection data that records what was actually built. What was previously a sequence of disconnected judgments is now a quantitatively connected, semi-automated workflow. The path from this foundation to a fully validated, industrially deployed smart process planning system is long and the future work identified in this chapter is substantial, but the foundation is in place, the tools are operational, and the results on the Quartic benchmark demonstrate that the approach is sound. The framework is ready to be tested against the physical world.

REFERENCES

[1] A. Moropoulou, A. Bakolas, and S. Anagnostopoulou, "Composite materials in ancient structures," *Cement and Concrete Composites*, Cement and Concrete Research in Greece, vol. 27, no. 2, pp. 295–300, Feb. 1, 2005, ISSN: 0958-9465. DOI: 10.1016/j.cemconcomp.2004.02.018 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0958946504000356

[2] T.-D. Ngo, *Composite and Nanocomposite Materials: From Knowledge to Industrial Applications*. BoD – Books on Demand, Jul. 15, 2020, 242 pp., ISBN: 978-1-78985-390-2. Google Books: 9kr9DwAAQBAJ.

[3] G. Lubin, *Handbook of Composites*. Springer US, Jun. 30, 1982, 966 pp., ISBN: 978-0-442-24897-0. Google Books: l79RAAAAMAAJ.

[4] D. Hull and T. W. Clyne, *An Introduction to Composite Materials*, 2nd ed. Cambridge University Press, Aug. 13, 1996, ISBN: 978-0-521-38190-1 978-0-521-38855-9 978-1-139-17013-0. DOI: 10.1017/CBO9781139170130 Accessed: Oct. 11, 2025. [Online]. Available: https://www.cambridge.org/core/product/identifier/9781139170130/type/book

[5] J. Hale, "Boeing 787: From the Ground Up," *Aero Magazine*, pp. 16–23, 2006. [Online]. Available: https://www.researchgate.net/publication/235947459_787_From_The_Ground_Up

[6] D. Evrard and F. Alonso, "FAST (Flight Airworthiness Support Technology) - Special Edition: A350 XWB," *Airbus Technical Magazine*, Jun. 2013, ISSN: 1293-5476. [Online]. Available: https://aircraft.airbus.com/sites/g/files/jlcbta126/files/2022-04/FAST_specialA350.pdf

[7] G. Rousseau, R. Wehbe, J. Halbritter, and R. Harik, "Automated Fiber Placement Path Planning: A state-of-the-art review," *Computer-Aided Design and Applications*, vol. 16, no. 2, pp. 172–203, Aug. 13, 2018, ISSN: 16864360. DOI: 10.14733/cadaps.2019.172-203 Accessed: Oct. 11, 2025. [Online]. Available: http://cad-journal.net/files/vol_16/Vol16No2.html

[8] A. Quilter, "Composites in Aerospace Applications," White Paper. [Online]. Available: https://www.spglobal.com/pdf/Composites-Aerospace-Applications-whitepaper_264558110913046532.pdf

[9] F. Khan et al., "Advances of composite materials in automobile applications – A review," *Journal of Engineering Research*, vol. 13, no. 2, pp. 1001–1023, Jun. 1, 2025, ISSN: 2307-1877. DOI: 10.1016/j.jer.2024.02.017 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2307187724000440

[10] T. Trzepieciński, T. Batu, F. Kibrete, and H. G. Lemu, "Application of Composite Materials for Energy Generation Devices," *Journal of Composites Science*, vol. 7, no. 2, p. 55, Feb. 3, 2023, ISSN: 2504-477X. DOI: 10.3390/jcs7020055 Accessed: Oct. 11, 2025. [Online]. Available: https://www.mdpi.com/2504-477X/7/2/55

[11] R. Harik and M. Godbold, "Towards Smart Automated Fiber Placement," in *ASM Handbook: Composites*, vol. 22, 22 vols., ASM International.

[12] J. Zhang, G. Lin, U. Vaidya, and H. Wang, "Past, present and future prospective of global carbon fibre composite developments and applications," *Composites Part B: Engineering*, vol. 250, p. 110 463, Feb. 1, 2023, ISSN: 1359-8368. DOI: 10.1016/j.compositesb.2022.110463 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1359836822008368

[13] H. Ghiasi, D. Pasini, and L. Lessard, "Optimum stacking sequence design of composite materials Part I: Constant stiffness design," *Composite Structures*, vol. 90, no. 1, pp. 1–11, Sep. 1, 2009, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2009.01.006 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S026382230900018X

[14] M. Sultana. "Micromechanics of Composites: Understanding the Building Blocks," Medium, Accessed: Oct. 11, 2025. [Online]. Available: https://masrath-sultana.medium.com/micromechanics-of-composites-understanding-the-building-blocks-a432becf1c1

[15] B. E. Chitwood and M. S. Howeth, "Apparatus for making laminated structural shapes by the controlled detrusive placement and polymerization of tectonic filamentous tapes," U.S. Patent 3574040A, Apr. 6, 1971. Accessed: Oct. 18, 2025. [Online]. Available: https://patents.google.com/patent/US3574040A/en

[16] D. Tele, N. Wakhare, R. Bhosale, P. Bharde, and S. G. Nerkar, "A Review on Design and Development of Filament Winding Machine for Composite Materials," 2016. Accessed: Oct. 18, 2025. [Online]. Available: https://www.semanticscholar.org/paper/A-Review-on-Design-and-Development-of-Filament-for-Tele-Wakhare/fcc2e7a13a2363de68982fcab2d8521a83205cff

[17] E. Hardesty, W. Goldsworthy, and H. Karlson, "Geodesic path length compensator for composite-tape placement head," U.S. Patent 3810805A, May 14, 1974. Accessed: Oct. 11, 2025. [Online]. Available: https://patents.google.com/patent/US3810805A/en

[18] R. Harik and A. Brasington, *Automated Fiber Placement: Status, Challenges, and Evolution*. SAE International, Jun. 11, 2025, 278 pp., ISBN: 978-1-4686-0306-4. Google Books: 0npoEQAAQBAJ.

[19] A. Brasington, C. Sacco, J. Halbritter, R. Wehbe, and R. Harik, "Automated fiber placement: A review of history, current technologies, and future paths forward,"

Composites Part C: Open Access, vol. 6, p. 100 182, Oct. 1, 2021, ISSN: 2666-6820. DOI: 10.1016/j.jcomc.2021.100182 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666682021000773

[20] R. Grenoble, R. Harik, D. Munden, J. Halbritter, D. Jegley, and B. Mason, "Assessment of Automated Fiber Placement for the Fabrication of Composite Wind Tunnel Blades," presented at the SAMPE, Diamond Bar, CA, United States, May 20, 2019. Accessed: Oct. 18, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002664

[21] W. Huberty, M. Roberson, B. Cai, M. Hendrickson, and Mississippi State University: Advanced Composites Institute, "State of the Industry – Resin Infusion: A Literature Review," Federal Aviation Administration: William J. Hughes Technical Center, May 1, 2024. DOI: 10.21949/1528268 Accessed: Feb. 13, 2026. [Online]. Available: https://rosap.ntl.bts.gov/view/dot/75056

[22] A. R. Brasington, "Comprehensive Process Planning Optimization Framework for Automated Fiber Placement," Ph.D. dissertation, University of South Carolina, 2023. Accessed: Oct. 17, 2025. [Online]. Available: https://search.proquest.com/openview/79fe6c73ab188f3de8cb9bd7d0c2af57/1?pq-origsite=gscholar&cbl=18750&diss=y

[23] R. J. Cano, H. L. Belvin, A. B. Hulcher, and R. W. Grenoble, "Studies on Automated Manufacturing of High Performance Composites," presented at the Structure Specialist Meeting, Williamsburg, VA, Jan. 1, 2001. Accessed: Oct. 11, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20040086087

[24] R. Calawa and J. Nancarrow, "Medium Wave Infrared Heater for High-Speed Fiber Placement," presented at the Aerospace Technology Conference and Exposition, Los Angeles, California, United States, Sep. 17, 2007, pp. 2007-01–3842. DOI: 10.4271/2007-01-3842 Accessed: Oct. 11, 2025. [Online]. Available: https://saemobilus.sae.org/papers/medium-wave-infrared-heater-high-speed-fiber-placement-2007-01-3842

[25] J. Jiang, Y. He, H. Wang, and Y. Ke, "Modeling and experimental validation of compaction pressure distribution for automated fiber placement," *Composite Structures*, vol. 256, p. 113 101, Jan. 15, 2021, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2020.113101 Accessed: Oct. 11, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263822320330270

[26] A. Halbritter, "Big Parts Demand Big Changes to the Fiber Placement Status Quo," presented at the SME Composites Manufacturing (Mesa, AZ), 2012.

[27] T. Rudberg, J. Nielson, M. Henscheid, and J. Cemenska, "Improving AFP Cell Performance," *SAE International Journal of Aerospace*, vol. 07, no. 2, pp. 317–321, Sep. 16, 2014, ISSN: 1946-3855, 1946-3901. DOI: 10.4271/2014-01-2272 Accessed:

Oct. 11, 2025. [Online]. Available: https://saemobilus.sae.org/articles/improving-afp-cell-performance-2014-01-2272

[28] R. DeVlieg, K. Jeffries, and P. Vogeli, "High-Speed Fiber Placement on Large Complex Structures," presented at the Aerospace Technology Conference and Exposition, Los Angeles, California, United States, Sep. 17, 2007, pp. 2007-01–3843. DOI: 10.4271/2007-01-3843 Accessed: Oct. 11, 2025. [Online]. Available: https://saemobilus.sae.org/papers/high-speed-fiber-placement-large-complex-structures-2007-01-3843

[29] R. Elhajjar, "Production Implementation of Multiple Machine, High Speed Fiber Placement for Large Structures," in *Additive Manufacturing of Aerospace Composite Structures: Fabrication and Reliability*, SAE, 2017, pp. 7–14, ISBN: 978-0-7680-8406-1. Accessed: Oct. 11, 2025. [Online]. Available: https://ieeexplore.ieee.org/document/8504647

[30] K. A. Jeffries, "Enhanced Robotic Automated Fiber Placement with Accurate Robot Technology and Modular Fiber Placement Head," *SAE International Journal of Aerospace*, vol. 06, no. 2, pp. 774–779, Sep. 17, 2013, ISSN: 1946-3855, 1946-3901. DOI: 10.4271/2013-01-2290 Accessed: Oct. 11, 2025. [Online]. Available: https://saemobilus.sae.org/articles/enhanced-robotic-automated-fiber-placement-accurate-robot-technology-modular-fiber-placement-head-2013-01-2290

[31] I. Martin, D. Saenz del Castillo, A. Fernandez, and A. Güemes, "Advanced Thermoplastic Composite Manufacturing by In-Situ Consolidation: A Review," *Journal of Composites Science*, vol. 4, no. 4, p. 149, Dec. 2020, ISSN: 2504-477X. DOI: 10.3390/jcs4040149 Accessed: Oct. 11, 2025. [Online]. Available: https://www.mdpi.com/2504-477X/4/4/149

[32] J. G. Davis and H. L. Bohon, *First NASA Advanced Composites Technology Conference, Part 2*. Jan. 1, 1991. Accessed: Oct. 11, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/19930021652

[33] R. L. Anderson and C. G. Grant, "Advanced fiber placement of composite fuselage structures," Jan. 1, 1991. Accessed: Oct. 11, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/19930021675

[34] F. Bullock, S. Kowalski, and R. Young, "Automated prepreg tow placement for composite structures.," *ADVANCED MATERIALS: THE CHALLENGE FOR THE NEXT DECADE.*, vol. 1, pp. 734–745, 1990.

[35] C. G. Grant and V. M. Benson, "Automated fiber placement: Evolution and current demonstrations," Jan. 1, 1993. Accessed: Oct. 11, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/19950022411

[36] M. L. Enders and P. C. Hopkins, “Developments in the fiber placement process.,” in *36. INTERNATIONAL SAMPE SYMPOSIUM AND EXHIBITION, 1991,*, 1991, pp. 778–790.

[37] D. O. Evans, “Design considerations for fiber placement,” *Advanced Materials: Performance Through Technology Insertion.*, vol. 38, pp. 170–181, 1993.

[38] M. J. Pasanen, J. P. Martin, R. J. Langone, and J. A. Mondo, “Advanced composite fiber placement: Process to application,” in *PROCEEDINGS OF THE JAPAN INTERNATIONAL SAMPE SYMPOSIUM*, JAPAN CHAPTER OF SAMPE, 1997, pp. 1055–1060.

[39] R. Measom, “Fiber placement low-cost production for complex composite structures,” in *American Helicopter Society 52^< Nd> Annual Forum, 1996*, 1996.

[40] M. Gruber, M. Lamontia, and B. Waibel, “AUTOMATED FABRICATION PROCESSES FOR LARGE COMPOSITE AEROSPACE STRUCTURES : A TRADE STUDY,” 2007. Accessed: Oct. 11, 2025. [Online]. Available: https://www.semanticscholar.org/paper/AUTOMATED-FABRICATION-PROCESSES-FOR-LARGE-COMPOSITE-Gruber-Lamontia/077e64a168576d1afcd91cac25419cbca643fe15

[41] R. W. Engelbart, M. R. Chapman, B. A. Johnson, K. A. Soucy, R. Hannebaum, and S. Schrader, “Systems and methods enabling automated return to and/or repair of defects with a material placement machine,” pat., May 2, 2006. Accessed: Oct. 11, 2025. [Online]. Available: https://patents.google.com/patent/US7039485B2/en

[42] M. Kirkpatrick, “Digital Twins in Advanced Manufacturing to Enhance Manufacturing Efficiency,” M.S. thesis, University of South Carolina, Oct. 1, 2022. [Online]. Available: https://scholarcommons.sc.edu/etd/7092

[43] R. Harik, “neXt Automated Fiber Placement: Advancing Composites Manufacturing Towards a New Paradigm,” *SAMPE Journal*, pp. 6–14, 2020. Accessed: Oct. 17, 2025. [Online]. Available: https://www.nxtbook.com/nxtbooks/sampe/journal_20201112/index.php?startid=6#/p/6

[44] H. Shyu, *DoD INSTRUCTION 5000.97 DIGITAL ENGINEERING*, Department of Defense, Aug. 8, 2007. [Online]. Available: https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500097p.PDF?ver=bePIqKXaLUTK_Iu5iTNREw%3D%3D

[45] M. Hyer and H. Lee, “The use of curvilinear fiber format to improve buckling resistance of composite plates with central circular holes,” *Composite Structures*, vol. 18, no. 3, pp. 239–261, Jan. 1991, issn: 02638223. doi: 10.1016/0263-8223(91)90035-W Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/026382239190035W

[46] L. Esposito et al., “Topology optimization-guided stiffening of composites realized through Automated Fiber Placement,” *Composites Part B: Engineering*, vol. 164,

pp. 309–323, May 2019, ISSN: 13598368. DOI: 10.1016/j.compositesb.2018.11.032 Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836818330993

[47] Y. Blanchard, "Composites Design Optimization for Automated Fiber Placement Process," presented at the SAE 2014 Aerospace Manufacturing and Automated Fastening Conference & Exhibition, Salt Lake City, Utah, United States, Sep. 16, 2014, pp. 2014-01–2261. DOI: 10.4271/2014-01-2261 Accessed: Oct. 17, 2025. [Online]. Available: https://saemobilus.sae.org/papers/composites-design-optimization-automated-fiber-placement-process-2014-01-2261

[48] Y. D. Boon, S. C. Joshi, S. K. Bhudolia, and G. Gohel, "Recent Advances on the Design Automation for Performance-Optimized Fiber Reinforced Polymer Composite Components," *Journal of Composites Science*, vol. 4, no. 2, p. 61, May 29, 2020, ISSN: 2504-477X. DOI: 10.3390/jcs4020061 Accessed: Oct. 17, 2025. [Online]. Available: https://www.mdpi.com/2504-477X/4/2/61

[49] R. Harik, C. Saidy, S. J. Williams, Z. Gurdal, and B. Grimsley, "Automated Fiber Placement Defect Identity Cards: Cause, Anticipation, Existence, Significance, and Progression," presented at the SAMPE, Diamond Bar, CA, United States, May 21, 2018. Accessed: Oct. 15, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002536

[50] M. Godbold, "A Globalized Optimization Schema for Automated Fiber Placement Processing Parameters," M.S. thesis, University of South Carolina, Oct. 1, 2023. [Online]. Available: https://scholarcommons.sc.edu/etd/7644

[51] A. Brasington, B. Francis, M. Godbold, and R. Harik, "A review and framework for modeling methodologies to advance automated fiber placement," *Composites Part C: Open Access*, vol. 10, p. 100 347, Mar. 1, 2023, ISSN: 2666-6820. DOI: 10.1016/j.jcomc.2023.100347 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666682023000038

[52] J. Halbritter, R. Harik, C. Saidy, A. Noevere, and B. Grimsley, "Automation of AFP Process Planning Functions: Importance and Ranking," in *SAMPE 2019 - Charlotte, NC*, SAMPE, Apr. 11, 2019. DOI: 10.33599/nasampe/s.19.1592 Accessed: Oct. 12, 2025. [Online]. Available: https://sampe.knack.com/technical-papers#home/technicalpaperdetails/5cb7823071efac0cb6e45aa0/

[53] J. Halbritter, "Automation of Process Planning for Automated Fiber Placement," M.S. thesis, University of South Carolina, Apr. 1, 2020. [Online]. Available: https://scholarcommons.sc.edu/etd/5953

[54] C. S. Gillespie, "Integrating 20 Tow, 1.5 Inch AFP Heads on Large Part for High Deposition Rate Application," presented at the AeroTech Congress & Exhibition, Fort Worth, Texas, United States, Sep. 19, 2017, pp. 2017-01–2149. DOI: 10.4271/2017-01-2149 Accessed: Oct. 17, 2025. [Online]. Available: https://saemobilus.sae.org/

papers/integrating-20-tow-15-inch-afp-heads-large-part-high-deposition-rate-application-2017-01-2149

[55] V. Dhinakaran, K. Surendar, M. Hasunfur Riyaz, and M. Ravichandran, "Review on study of thermosetting and thermoplastic materials in the automated fiber placement process," *Materials Today: Proceedings*, vol. 27, pp. 812–815, 2020, ISSN: 22147853. DOI: 10.1016/j.matpr.2019.12.355 Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2214785319343858

[56] N. Bakhshi and M. Hojjati, "Effect of compaction roller on layup quality and defects formation in automated fiber placement," *Journal of Reinforced Plastics and Composites*, vol. 39, no. 1–2, pp. 3–20, Jan. 2020, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/0731684419868845 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0731684419868845

[57] J. Brüning, B. Denkena, M.-A. Dittrich, and T. Hocke, "Machine Learning Approach for Optimization of Automated Fiber Placement Processes," *Procedia CIRP*, vol. 66, pp. 74–78, 2017, ISSN: 22128271. DOI: 10.1016/j.procir.2017.03.295 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2212827117304833

[58] T. Rudberg, J. Cemenska, and E. Sherrard, "A Process for Delivering Extreme AFP Head Reliability," *SAE International Journal of Advances and Current Practices in Mobility*, vol. 1, no. 2, pp. 333–342, Mar. 19, 2019, ISSN: 2641-9645. DOI: 10.4271/2019-01-1349 Accessed: Oct. 16, 2025. [Online]. Available: https://www.sae.org/content/2019-01-1349

[59] J. Cemenska, T. Rudberg, and M. Henscheid, "Automated In-Process Inspection System for AFP Machines," *SAE International Journal of Aerospace*, vol. 08, no. 2, pp. 303–309, Sep. 15, 2015, ISSN: 1946-3855, 1946-3901. DOI: 10.4271/2015-01-2608 Accessed: Oct. 16, 2025. [Online]. Available: https://saemobilus.sae.org/articles/automated-process-inspection-system-afp-machines-2015-01-2608

[60] C. Sacco et al., "On the effect of manual rework in AFP quality control for a doubly-curved part," *Composites Part B: Engineering*, vol. 227, p. 109 432, Dec. 2021, ISSN: 13598368. DOI: 10.1016/j.compositesb.2021.109432 Accessed: Oct. 24, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836821008015

[61] P. D. Juarez, K. E. Cramer, and J. P. Seebo, "Advances in in situ inspection of automated fiber placement systems," presented at the SPIE Commercial + Scientific Sensing and Imaging, J. N. Zalameda and P. Bison, Eds., Baltimore, Maryland, United States, May 11, 2016, p. 986 109. DOI: 10.1117/12.2223028 Accessed: Oct. 17, 2025. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2223028

[62] B. Denkena, C. Schmidt, K. Völtzer, and T. Hocke, "Thermographic online monitoring system for Automated Fiber Placement processes," *Composites Part B: Engineering*, vol. 97, pp. 239–243, Jul. 2016, ISSN: 13598368. DOI: 10.1016/j.compositesb.2016.04.076 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836816305285

[63] E. D. Gregory and P. D. Juarez, "In-situ thermography of automated fiber placement parts," presented at the 44TH ANNUAL REVIEW OF PROGRESS IN QUANTITATIVE NONDESTRUCTIVE EVALUATION, VOLUME 37, Provo, Utah, USA, 2018, p. 060 005. DOI: 10.1063/1.5031551 Accessed: Oct. 17, 2025. [Online]. Available: https://pubs.aip.org/aip/acp/article/819573

[64] C. Schmidt, B. Denkena, K. Völtzer, and T. Hocke, "Thermal Image-based Monitoring for the Automated Fiber Placement Process," *Procedia CIRP*, vol. 62, pp. 27–32, 2017, ISSN: 22128271. DOI: 10.1016/j.procir.2016.06.058 Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S221282711630676X

[65] P. D. Juarez and E. D. Gregory, "In situ thermal inspection of automated fiber placement manufacturing," presented at the 45TH ANNUAL REVIEW OF PROGRESS IN QUANTITATIVE NONDESTRUCTIVE EVALUATION, VOLUME 38, Vermont, USA, 2019, p. 120 005. DOI: 10.1063/1.5099847 Accessed: Oct. 17, 2025. [Online]. Available: https://pubs.aip.org/aip/acp/article/888679

[66] E. Hu and F. Haifeng, "Surface profile inspection of a moving object by using dual-frequency Fourier transform profilometry," *Optik*, vol. 122, no. 14, pp. 1245–1248, Jul. 2011, ISSN: 00304026. DOI: 10.1016/j.ijleo.2010.08.007 Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0030402610003773

[67] C. Sacco, "Machine Learning Methods for Rapid Inspection of Automated Fiber Placement Manufactured Composite Structures," M.S. thesis, University of South Carolina, Jul. 1, 2019. [Online]. Available: https://scholarcommons.sc.edu/etd/5475

[68] C. Sacco, A. B. Radwan, R. Harik, and M. Van Tooren, "Automated Fiber Placement Defects: Automated Inspection and Characterization," presented at the Society for the Advancement of Materials and Process Engineering, Covina, CA, United States, May 21, 2018. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20190027133

[69] C. Schmidt, C. Schultz, P. Weber, and B. Denkena, "Evaluation of eddy current testing for quality assurance and process monitoring of automated fiber placement," *Composites Part B: Engineering*, vol. 56, pp. 109–116, Jan. 2014, ISSN: 13598368. DOI: 10.1016/j.compositesb.2013.08.061 Accessed: Oct. 17, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836813004691

[70] H. Heuer, M. Schulze, and N. Meyendorf, "High Resolution Inspection of Carbon Fiber Materials by Eddy Current Techniques," presented at the 2nd Int. Symposium

on NDT in Aerospace, vol. 16, Mar. 1, 2011. Accessed: Oct. 17, 2025. [Online]. Available: https://www.ndt.net//search/docs.php3

[71] C. Sacco, A. Baz Radwan, A. Anderson, R. Harik, and E. Gregory, "Machine learning in composites manufacturing: A case study of Automated Fiber Placement inspection," *Composite Structures*, vol. 250, p. 112 514, Oct. 2020, ISSN: 02638223. DOI: 10.1016/j.compstruct.2020.112514 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822320313659

[72] C. X. Li et al., "Manufacturing Autoclave-Grade Thermoset Carbon Fiber-Reinforced Polymer Aerospace Composites without an Autoclave Using Nanoporous Materials," *ACS Applied Materials & Interfaces*, vol. 16, no. 19, pp. 25 280–25 293, May 15, 2024, ISSN: 1944-8244, 1944-8252. DOI: 10.1021/acsami.4c02107 Accessed: Oct. 17, 2025. [Online]. Available: https://pubs.acs.org/doi/10.1021/acsami.4c02107

[73] K. Yassin and M. Hojjati, "Processing of thermoplastic matrix composites through automated fiber placement and tape laying methods: A review," *Journal of Thermoplastic Composite Materials*, vol. 31, no. 12, pp. 1676–1725, Dec. 2018, ISSN: 0892-7057, 1530-7980. DOI: 10.1177/0892705717738305 Accessed: Oct. 17, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0892705717738305

[74] M. Belhaj et al., "Dry fiber automated placement of carbon fibrous preforms," *Composites Part B: Engineering*, vol. 50, pp. 107–111, Jul. 2013, ISSN: 13598368. DOI: 10.1016/j.compositesb.2013.01.014 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836813000206

[75] A. Brasington, C. Smith, J. Halbritter, R. Wehbe, and R. Harik, "Surrogate based methods for rapid starting point optimization in automated fiber placement," in *SAMPE 2022 Conference Proceedings*, 2022.

[76] N. Swingle, "Automated Fiber Placement Through Thickness Defect Stacking Optimization," M.S. thesis, University of South Carolina, Jul. 1, 2023. [Online]. Available: https://scholarcommons.sc.edu/etd/7493

[77] N. Patel, "AUTOMATED FIBER PLACEMENT LAMINATE LEVEL OPTIMIZATION: A PHYSICSINSPIRED METHOD TO ANALYZE THROUGH-THICKNESS DEFECT INTERACTIONS," M.S. thesis, University of South Carolina, 2025.

[78] N. Scheirer, S. D. Holland, and A. Krishnamurthy, "Fiber Layup Generation on Curved Composite Structures," *Computer-Aided Design*, vol. 136, p. 103 031, Jul. 1, 2021, ISSN: 0010-4485. DOI: 10.1016/j.cad.2021.103031 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010448521000427

[79] F. Yang et al., "An efficient pre-analysis and optimization generation method for reference curves of automated fiber placement path planning," *Journal of Composite Materials*, vol. 58, no. 26, pp. 2781–2797, Nov. 1, 2024, ISSN: 0021-9983. DOI:

10.1177/00219983241270958 Accessed: Oct. 12, 2025. [Online]. Available: https://doi.org/10.1177/00219983241270958

[80] N. M. Patrikalakis and T. Maekawa, *Shape Interrogation for Computer Aided Design and Manufacturing*. Berlin, Heidelberg: Springer, 2002, ISBN: 978-3-642-04073-3 978-3-642-04074-0. DOI: 10.1007/978-3-642-04074-0 Accessed: Oct. 12, 2025. [Online]. Available: https://link.springer.com/10.1007/978-3-642-04074-0

[81] Z. Gurdal, B. Tatting, and K. Wu, "Tow-Placement Technology and Fabrication Issues for Laminated Composite Structures," in *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Austin, Texas: American Institute of Aeronautics and Astronautics, Apr. 18, 2005, ISBN: 978-1-62410-065-9. DOI: 10.2514/6.2005-2017 Accessed: Oct. 12, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2005-2017

[82] P. Pourabdollah et al., "Automated composite design generation and optimization using variable angle tow plies for aerospace applications," *Composite Structures*, vol. 372, p. 119 511, Nov. 15, 2025, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2025.119511 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263822325006762

[83] M. A. Albazzan, R. Harik, B. F. Tatting, and Z. Gürdal, "Efficient design optimization of nonconventional laminated composites using lamination parameters: A state of the art," *Composite Structures*, vol. 209, pp. 362–374, Feb. 1, 2019, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2018.10.095 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263822318321664

[84] M. Albazzan, "Efficient Design Optimization Methodology for Manufacturable Variable Stiffness Laminated Composite Structures," *Theses and Dissertations*, Apr. 1, 2020. [Online]. Available: https://scholarcommons.sc.edu/etd/5731

[85] Z. Gürdal, B. F. Tatting, and C. K. Wu, "Variable stiffness composite panels: Effects of stiffness variation on the in-plane and buckling response," *Composites Part A: Applied Science and Manufacturing*, vol. 39, no. 5, pp. 911–922, May 1, 2008, ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2007.11.015 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1359835X07002643

[86] K. Schueler, J. Miller, and R. Hale, "Approximate Geometric Methods in Application to the Modeling of Fiber Placed Composite Structures," *Journal of Computing and Information Science in Engineering*, vol. 4, no. 3, pp. 251–256, Sep. 1, 2004, ISSN: 1530-9827, 1944-7078. DOI: 10.1115/1.1736685 Accessed: Oct. 12, 2025. [Online]. Available: https://asmedigitalcollection.asme.org/computingengineering/article/4/3/251/445337/Approximate-Geometric-Methods-in-Application-to

[87] B. Shirinzadeh, G. Cassidy, D. Oetomo, G. Alici, and M. H. Ang Jr, "Trajectory generation for open-contoured structures in robotic fibre placement," *Robotics and Computer-Integrated Manufacturing*, vol. 23, no. 4, pp. 380–394, Aug. 1, 2007, ISSN:

0736-5845. doi: 10.1016/j.rcim.2006.04.006 Accessed: Oct. 12, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0736584506000408

[88] R. Wehbe, "Modeling Of Tow Wrinkling In Automated Fiber Placement Based On Geometrical Considerations," M.S. thesis, University of South Carolina, Jan. 1, 2017. [Online]. Available: https://scholarcommons.sc.edu/etd/4449

[89] X. Wang, W. Zhang, and L. Zhang, "Intersection of a ruled surface with a free-form surface," *Numerical Algorithms*, vol. 46, no. 1, pp. 85–100, Oct. 30, 2007, issn: 1017-1398, 1572-9265. doi: 10.1007/s11075-007-9118-y Accessed: Oct. 16, 2025. [Online]. Available: http://link.springer.com/10.1007/s11075-007-9118-y

[90] L. Yan, Z. C. Chen, Y. Shi, and R. Mo, "An accurate approach to roller path generation for robotic fibre placement of free-form surface composites," *Robotics and Computer-Integrated Manufacturing*, vol. 30, no. 3, pp. 277–286, Jun. 2014, issn: 07365845. doi: 10.1016/j.rcim.2013.10.007 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0736584513000884

[91] M. Bruyneel and S. Zein, "A modified Fast Marching Method for defining fiber placement trajectories over meshes," *Computers & Structures*, vol. 125, pp. 45–52, Sep. 2013, issn: 00457949. doi: 10.1016/j.compstruc.2013.04.015 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0045794913001454

[92] E. Lemaire, S. Zein, and M. Bruyneel, "Optimization of composite structures with curved fiber trajectories," *Composite Structures*, vol. 131, pp. 895–904, Nov. 2015, issn: 02638223. doi: 10.1016/j.compstruct.2015.06.040 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822315005048

[93] A. W. Blom, P. B. Stickler, and Z. Gürdal, "Optimization of a composite cylinder under bending by tailoring stiffness properties in circumferential direction," *Composites Part B: Engineering*, vol. 41, no. 2, pp. 157–165, Mar. 2010, issn: 13598368. doi: 10.1016/j.compositesb.2009.10.004 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836809001887

[94] A. W. Blom, C. S. Lopes, P. J. Kromwijk, Z. Gurdal, and P. Camanho, "A Theoretical Model to Study the Influence of Tow-drop Areas on the Stiffness and Strength of Variable-stiffness Laminates," *Journal of Composite Materials*, vol. 43, no. 5, pp. 403–425, Mar. 2009, issn: 0021-9983, 1530-793X. doi: 10.1177/0021998308097675 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0021998308097675

[95] A. W. Blom, S. Setoodeh, J. M. Hol, and Z. Gürdal, "Design of variable-stiffness conical shells for maximum fundamental eigenfrequency," *Computers & Structures*, vol. 86, no. 9, pp. 870–878, May 2008, issn: 00457949. doi: 10.1016/j.compstruc.

2007.04.020 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0045794907001630

[96] D. Jegley, B. Tatting, and Z. Gurdal, "Optimization of Elastically Tailored Tow-Placed Plates with Holes," in *44th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Norfolk, Virginia: American Institute of Aeronautics and Astronautics, Apr. 7, 2003, ISBN: 978-1-62410-100-7. DOI: 10.2514/6.2003-1420 Accessed: Oct. 16, 2025. [Online]. Available: http://arc.aiaa.org/doi/10.2514/6.2003-1420

[97] B. Kim, K. Hazra, P. Weaver, and K. Potter, "Limitations of fibre placement techniques for variable angle tow composites and their process-induced defects," Aug. 1, 2011. Accessed: Oct. 16, 2025. [Online]. Available: https://www.semanticscholar.org/paper/Limitations-of-fibre-placement-techniques-for-angle-Kim-Hazra/31322f6ac16aedfead2c2fba7c754d84149c2c42

[98] C. Lopes, P. Camanho, Z. Gürdal, and B. Tatting, "Progressive failure analysis of tow-placed, variable-stiffness composite panels," *International Journal of Solids and Structures*, vol. 44, no. 25–26, pp. 8493–8516, Dec. 2007, ISSN: 00207683. DOI: 10.1016/j.ijsolstr.2007.06.029 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0020768307002624

[99] C. Lopes, Z. Gürdal, and P. Camanho, "Variable-stiffness composite panels: Buckling and first-ply failure improvements over straight-fibre laminates," *Computers & Structures*, vol. 86, no. 9, pp. 897–907, May 2008, ISSN: 00457949. DOI: 10.1016/j.compstruc.2007.04.016 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0045794907001654

[100] B. F. Tatting, Z. Guerdal, and D. Jegley, "Design and Manufacture of Elastically Tailored Tow Placed Plates," NASA/CR-2002-211919, Aug. 1, 2002. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20020073162

[101] K. Wu and Z. Gurdal, "Thermal testing of tow-placed, variable stiffness panels," in *19th AIAA Applied Aerodynamics Conference*, Anaheim,CA,U.S.A.: American Institute of Aeronautics and Astronautics, Jun. 11, 2001. DOI: 10.2514/6.2001-1190 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2001-1190

[102] C. Wu, Z. Gurdal, and J. Starnes, "Structural Response of Compression-Loaded, Tow-Placed, Variable Stiffness Panels," in *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Denver, Colorado: American Institute of Aeronautics and Astronautics, Apr. 22, 2002, ISBN: 978-1-62410-117-5. DOI: 10.2514/6.2002-1512 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2002-1512

[103] M. Favaloro and D. Hauber, "Process and design considerations for the automated fiber placement process," presented at the SAMPE, Cincinnati, OH, 2007.

[104] E. Oromiehie, G. Prusty, P. Compston, and G. Rajan, "The influence of consolidation force on the performance of AFP manufacturedThe influence of consolidation force on the performance of AFP manufactured laminateslaminates," presented at the 21st International Conference on Composite Materials (ICCM-21), 2017, pp. 1–11. [Online]. Available: https://ro.uow.edu.au/eispapers1/1810

[105] M. A. Khan, P. Mitschang, and R. Schledjewski, "Parametric study on processing parameters and resulting part quality through thermoplastic tape placement process," *Journal of Composite Materials*, vol. 47, no. 4, pp. 485–499, Feb. 2013, ISSN: 0021-9983, 1530-793X. DOI: 10.1177/0021998312441810 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0021998312441810

[106] D. Felderhoff and V. Steiner, "A New Compact Robotic Head for Thermoplastic Fiber Placement.," *International SAMPE Symposium and Exhibition*, vol. 38, no. 1, pp. 138–151, 1993. Accessed: Oct. 15, 2025. [Online]. Available: https://jglobal.jst.go.jp/en/detail?JGLOBAL_ID=200902189958781145

[107] D. Saenz del Castillo, I. Martín, F. Rodriguez-Lence, and A. Güemes, "On-line monitoring of a laser assisted fiber placement process with CFR thermoplastic matrix by using Fiber Bragg Gratings," *e-Journal of Nondestructive Testing*, vol. 21, no. 08, Aug. 1, 2016. Accessed: Oct. 15, 2025. [Online]. Available: https://www.ndt.net//search/docs.php3

[108] T. T. Zacchia, F. Shadmehri, J. Fortin-Simpson, and S. V. Hoa, "E: Design Of Hard Compaction Rollers For Automated Fiber Placement On Complex Mandrel Geometries," in *Progress in Canadian Mechanical Engineering*, York University Libraries, May 30, 2018, ISBN: 978-1-77355-023-7. DOI: 10.25071/10315/35245 HDL: 10315/35245. Accessed: Oct. 15, 2025. [Online]. Available: http://hdl.handle.net/10315/35245

[109] C. Schmidt, B. Denkena, T. Hocke, and K. Völtzer, "Influence of AFP Process Parameters on the Temperature Distribution Used for Thermal in-process Monitoring," *Procedia CIRP*, vol. 66, pp. 68–73, 2017, ISSN: 22128271. DOI: 10.1016/j.procir.2017.03.220 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2212827117304031

[110] R. Y. Wehbe, R. Harik, and Z. Gurdal, "In-plane tow deformations due to steering in automated fiber placement," in *AIAA Scitech 2019 Forum*, San Diego, California: American Institute of Aeronautics and Astronautics, Jan. 7, 2019, ISBN: 978-1-62410-578-4. DOI: 10.2514/6.2019-1271 Accessed: Oct. 15, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2019-1271

[111] R. Lichtinger, J. Lacalle, R. Hinterhölzl, U. Beier, and K. Drechsler, "Simulation and experimental validation of gaps and bridging in the automated fiber placement process," *Science and Engineering of Composite Materials*, vol. 22, no. 2, pp. 131–148, Mar. 1, 2015, ISSN: 2191-0359, 0792-1233. DOI: 10.1515/secm-2013-0158

Accessed: Oct. 15, 2025. [Online]. Available: https://www.degruyter.com/document/doi/10.1515/secm-2013-0158/html

[112] J. Cheng, D. Zhao, K. Liu, Y. Wang, and H. Chen, “Modeling and impact analysis on contact characteristic of the compaction roller for composite automated placement,” *Journal of Reinforced Plastics and Composites*, vol. 37, no. 23, pp. 1418–1432, Dec. 2018, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/0731684418798151 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0731684418798151

[113] W. Roudy, Z. Kevin, M. William, C. Sacco, and R. Harik, “Investigation of the Influence of Roller Compaction on Quality of Defects in Automated Fiber Placement,” presented at the CAMX, 2021.

[114] P. Schaefer, D. Gierszewski, A. Kollmannsberger, S. Zaremba, and K. Drechsler, “Analysis and improved process response prediction of laser- assisted automated tape placement with PA-6/carbon tapes using Design of Experiments and numerical simulations,” *Composites Part A: Applied Science and Manufacturing*, vol. 96, pp. 137–146, May 2017, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2017.02.008 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X17300568

[115] K. J. Ahn, J. C. Seferis, T. Pelton, and M. Wilhelm, “Analysis and characterization of prepreg tack,” *Polymer Composites*, vol. 13, no. 3, pp. 197–206, Jun. 1992, ISSN: 0272-8397, 1548-0569. DOI: 10.1002/pc.750130308 Accessed: Oct. 15, 2025. [Online]. Available: https://4spepublications.onlinelibrary.wiley.com/doi/10.1002/pc.750130308

[116] R. Crossley, P. Schubel, and N. Warrior, “The experimental determination of prepreg tack and dynamic stiffness,” *Composites Part A: Applied Science and Manufacturing*, vol. 43, no. 3, pp. 423–434, Mar. 2012, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2011.10.014 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X11003538

[117] O. Dubois, J.-B. Le Cam, and A. Béakou, “Experimental Analysis of Prepreg Tack,” *Experimental Mechanics*, vol. 50, no. 5, pp. 599–606, Jun. 2010, ISSN: 0014-4851, 1741-2765. DOI: 10.1007/s11340-009-9236-7 Accessed: Oct. 15, 2025. [Online]. Available: http://link.springer.com/10.1007/s11340-009-9236-7

[118] R. Crossley, P. Schubel, and D. De Focatiis, “Time–temperature equivalence in the tack and dynamic stiffness of polymer prepreg and its application to automated composites manufacturing,” *Composites Part A: Applied Science and Manufacturing*, vol. 52, pp. 126–133, Sep. 2013, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2013.05.002 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X1300136X

[119] A. Endruweit et al., "Characterisation of tack for uni-directional prepreg tape employing a continuous application-and-peel test method," *Composites Part A: Applied Science and Manufacturing*, vol. 114, pp. 295–306, Nov. 2018, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2018.08.027 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X18303427

[120] A. Endruweit, D. S. De Focatiis, N. A. Warrior, S. Ghose, B. Johnson, and D. R. Younkin, "Characterization of Prepreg Tack to Aid Automated Material Placement," presented at the SAMPE, 2016. Accessed: Oct. 15, 2025. [Online]. Available: https://www.nasampe.org/store/viewproduct.aspx?id=6789393

[121] P.-E. Bourban, N. Bernet, J.-E. Zanetto, and J.-A. E. Månson, "Material phenomena controlling rapid processing of thermoplastic composites," *Composites Part A: Applied Science and Manufacturing*, vol. 32, no. 8, pp. 1045–1057, Aug. 2001, ISSN: 1359835X. DOI: 10.1016/S1359-835X(01)00017-3 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X01000173

[122] T. Gutowski, *Advanced Composites Manufacturing*. 1997, 600 pp., ISBN: 978-0-471-15301-6. Accessed: Oct. 15, 2025. [Online]. Available: https://www.wiley.com/en-us/Advanced+Composites+Manufacturing-p-9780471153016

[123] F. O. Sonmez and H. T. Hahn, "Analysis of the On-Line Consolidation Process in Thermoplastic Composite Tape Placement," *Journal of Thermoplastic Composite Materials*, vol. 10, no. 6, pp. 543–572, Nov. 1997, ISSN: 0892-7057, 1530-7980. DOI: 10.1177/089270579701000604 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/089270579701000604

[124] F. O. Sonmez and H. T. Hahn, "Modeling of Heat Transfer and Crystallization in Thermoplastic Composite Tape Placement Process," *Journal of Thermoplastic Composite Materials*, vol. 10, no. 3, pp. 198–240, May 1997, ISSN: 0892-7057, 1530-7980. DOI: 10.1177/089270579701000301 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/089270579701000301

[125] F. O. Sonmez, H. T. Hahn, and M. Akbulut, "Analysis of Process-Induced Residual Stresses in Tape Placement," *Journal of Thermoplastic Composite Materials*, vol. 15, no. 6, pp. 525–544, Nov. 2002, ISSN: 0892-7057, 1530-7980. DOI: 10.1177/0892705702015006207 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0892705702015006207

[126] M. A. Khan, P. Mitschang, and R. Schledjewski, "Tracing the Void Content Development and Identification of its Effecting Parameters during in Situ Consolidation of Thermoplastic Tape Material," *Polymers and Polymer Composites*, vol. 18, no. 1, pp. 1–15, Jan. 2010, ISSN: 0967-3911, 1478-2391. DOI: 10.1177/096739111001800101 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/096739111001800101

[127] S. Ranganathan, S. G. Advani, and M. A. Lamontia, “A Non-Isothermal Process Model for Consolidation and Void Reduction during In-Situ Tow Placement of Thermoplastic Composites,” *Journal of Composite Materials*, vol. 29, no. 8, pp. 1040–1062, May 1995, ISSN: 0021-9983, 1530-793X. DOI: 10.1177/002199839502900803 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/002199839502900803

[128] F. Yang and R. Pitchumani, “Nonisothermal healing and interlaminar bond strength evolution during thermoplastic matrix composites processing,” *Polymer Composites*, vol. 24, no. 2, pp. 263–278, Apr. 2003, ISSN: 0272-8397, 1548-0569. DOI: 10.1002/pc.10027 Accessed: Oct. 15, 2025. [Online]. Available: https://4spepublications.onlinelibrary.wiley.com/doi/10.1002/pc.10027

[129] S. C. Mantell and G. S. Springer, “Manufacturing Process Models for Thermoplastic Composites,” *Journal of Composite Materials*, vol. 26, no. 16, pp. 2348–2377, Jan. 1992, ISSN: 0021-9983, 1530-793X. DOI: 10.1177/002199839202601602 Accessed: Oct. 15, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/002199839202601602

[130] K. Kozaczuk, “AUTOMATED FIBER PLACEMENT SYSTEMS OVERVIEW,” *Transactions of the Institute of Aviation*, vol. 245, no. 4, pp. 52–59, Dec. 30, 2016, ISSN: 0509-6669, 2300-5408. DOI: 10.5604/05096669.1226355 Accessed: Oct. 15, 2025. [Online]. Available: http://9807.indexcopernicus.com/abstracted.php?level=5&ICID=1226355

[131] R. Engelhardt, R. Irmanputra, K. Brath, N. Aufenanger, and K. Drechsler, “Thermoset Prepreg Compaction during Automated Fiber Placement and Vacuum Debulking,” *Procedia CIRP*, vol. 85, pp. 153–158, 2019, ISSN: 22128271. DOI: 10.1016/j.procir.2019.09.025 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2212827119312855

[132] J. Chen, T. Chen-Keat, M. Hojjati, A. Vallee, M.-A. Octeau, and A. Yousefpour, “Impact of layup rate on the quality of fiber steering/cut-restart in automated fiber placement processes,” *Science and Engineering of Composite Materials*, vol. 22, no. 2, pp. 165–173, Mar. 1, 2015, ISSN: 2191-0359, 0792-1233. DOI: 10.1515/secm-2013-0257 Accessed: Oct. 15, 2025. [Online]. Available: https://www.degruyter.com/document/doi/10.1515/secm-2013-0257/html

[133] M. A. Khan, P. Mitschang, and R. Schledjewski, “Identification of some optimal parameters to achieve higher laminate quality through tape placement process,” *Advances in Polymer Technology*, vol. 29, no. 2, pp. 98–111, Jun. 2010, ISSN: 0730-6679, 1098-2329. DOI: 10.1002/adv.20177 Accessed: Oct. 15, 2025. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/adv.20177

[134] W. Zhang, F. Liu, Y. Lv, and X. Ding, “Modelling and layout design for an automated fibre placement mechanism,” *Mechanism and Machine Theory*, vol. 144, p. 103 651,

Feb. 2020, ISSN: 0094114X. DOI: 10.1016/j.mechmachtheory.2019.103651 Accessed: Oct. 15, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0094114X1931626X

[135] M. Assadi and T. Field, "AFP Processing of Dry Fiber Carbon Materials (DFP) for Improved Rates and Reliability," *SAE International Journal of Advances and Current Practices in Mobility*, vol. 2, no. 3, pp. 1196–1201, Mar. 10, 2020, ISSN: 2641-9645. DOI: 10.4271/2020-01-0030 Accessed: Oct. 15, 2025. [Online]. Available: https://www.sae.org/content/2020-01-0030

[136] K. Croft, L. Lessard, D. Pasini, M. Hojjati, J. Chen, and A. Yousefpour, "Experimental study of the effect of automated fiber placement induced defects on performance of composite laminates," *Composites Part A: Applied Science and Manufacturing*, vol. 42, no. 5, pp. 484–491, May 2011, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2011.01.007 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X11000224

[137] R. Anay et al., "An Experimental Investigation Concerning the Effects of AFP Defects on Progressive Failure of Tensile Coupons," in *AIAA Scitech 2019 Forum*, San Diego, California: American Institute of Aeronautics and Astronautics, Jan. 7, 2019, ISBN: 978-1-62410-578-4. DOI: 10.2514/6.2019-1547 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2019-1547

[138] A. Brasington, T. Schachner, and R. Harik, "3D Modeling and Printing of Automated Fiber Placement Defects," presented at the CAMX, 2020. [Online]. Available: https://www.nasampe.org/store/viewproduct.aspx?id=17557479

[139] A. Marouene, P. Legay, and R. Boukhili, "Experimental and numerical investigation on the open-hole compressive strength of AFP composites containing gaps and overlaps," *Journal of Composite Materials*, vol. 51, no. 26, pp. 3631–3646, Nov. 2017, ISSN: 0021-9983, 1530-793X. DOI: 10.1177/0021998317690917 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0021998317690917

[140] W. Woigk, S. R. Hallett, M. I. Jones, M. Kuhtz, A. Hornig, and M. Gude, "Experimental investigation of the effect of defects in Automated Fibre Placement produced composite laminates," *Composite Structures*, vol. 201, pp. 1004–1017, Oct. 2018, ISSN: 02638223. DOI: 10.1016/j.compstruct.2018.06.078 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822317339946

[141] M. Lan, D. Cartié, P. Davies, and C. Baley, "Influence of embedded gap and overlap fiber placement defects on the microstructure and shear and compression properties of carbon–epoxy laminates," *Composites Part A: Applied Science and Manufacturing*, vol. 82, pp. 198–207, Mar. 2016, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2015.12.007 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X15004716

[142] D. Maass, “Progress in automated ply inspection of AFP layups,” *Reinforced Plastics*, vol. 59, no. 5, pp. 242–245, Sep. 2015, ISSN: 00343617. DOI: 10.1016/j.repl.2015.05.002 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0034361715004002

[143] C. Sacco, A. B. Radwan, T. Beatty, and R. Harik, “Machine Learning Based AFP Inspection: A Tool for Characterization and Integration,” presented at the SAMPE, Diamond Bar, CA, United States, May 20, 2019. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002684

[144] M. H. Nguyen, P. Davidson, and A. M. Waas, “Experimental and numerical study on the tensile failure behavior of toughened-interlayer composite laminates with automated fiber placement (AFP) induced gap and overlap defects,” *International Journal of Material Forming*, vol. 14, no. 1, pp. 105–119, Jan. 2021, ISSN: 1960-6206, 1960-6214. DOI: 10.1007/s12289-020-01581-w Accessed: Feb. 19, 2026. [Online]. Available: https://link.springer.com/10.1007/s12289-020-01581-w

[145] L. Raps, A. R. Chadwick, and H. F. Voggenreiter, “Effect of geometry-induced triangular gap and overlap defects on in-situ AFP-manufactured structures,” *Composite Structures*, vol. 373, p. 119 556, Dec. 1, 2025, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2025.119556 Accessed: Oct. 16, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263822325007214

[146] S. Qian, X. Liu, Y. Ye, Q. Xu, T. Zhang, and X. Li, “Effect of gap and overlap fiber placement defects on the delamination behavior of L-shaped composite laminates,” *Composite Structures*, vol. 268, p. 113 963, Jul. 15, 2021, ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2021.113963 Accessed: Oct. 16, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263822321004232

[147] A. Sawicki and P. Minguett, “The effect of intraply overlaps and gaps upon the compression strength of composite laminates,” in *39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit*, Long Beach,CA,U.S.A.: American Institute of Aeronautics and Astronautics, Apr. 20, 1998. DOI: 10.2514/6.1998-1786 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.1998-1786

[148] Z. G. Ghauch, “Comprehensive Surrogate-Based Optimization of Lightweight Composite Manufacturing,” *AIAA Journal*, vol. 60, no. 7, pp. 4352–4366, Jul. 2022, ISSN: 0001-1452, 1533-385X. DOI: 10.2514/1.J061075 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/1.J061075

[149] C. Wanigasekara, E. Oromiehie, A. Swain, B. G. Prusty, and S. K. Nguang, “Machine Learning Based Predictive Model for AFP-Based Unidirectional Composite Laminates,” *IEEE Transactions on Industrial Informatics*, vol. 16, no. 4, pp. 2315–2324, Apr. 2020, ISSN: 1941-0050. DOI: 10.1109/TII.2019.2932398 Accessed: Oct. 16, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8784264

[150] M. Kheradpisheh and M. Hojjati, "A novel compaction roller with variable pressure distribution and contact time for automated fiber placement: Experimental and numerical analysis," *Composites Part A: Applied Science and Manufacturing*, vol. 190, p. 108 684, Mar. 2025, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2024.108684 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X24006821

[151] S. Meister, M. A. M. Wermes, J. Stüve, and R. M. Groves, "Review of image segmentation techniques for layup defect detection in the Automated Fiber Placement process: A comprehensive study to improve AFP inspection," *Journal of Intelligent Manufacturing*, vol. 32, no. 8, pp. 2099–2119, Dec. 2021, ISSN: 0956-5515, 1572-8145. DOI: 10.1007/s10845-021-01774-3 Accessed: Oct. 16, 2025. [Online]. Available: https://link.springer.com/10.1007/s10845-021-01774-3

[152] T. Wei, W. Qu, D. Yang, and Y. Ke, "Understanding and controlling wrinkle formation under time–process interactions during automated fiber placement (AFP) steering," *Composites Part A: Applied Science and Manufacturing*, vol. 199, p. 109 214, Dec. 1, 2025, ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2025.109214 Accessed: Oct. 16, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1359835X25005081

[153] M. Kheradpisheh and M. Hojjati, "Wrinkle Formation and Initial Defect Sensitivity of Steered Tow in Automated Fiber Placement," *Journal of Composites Science*, vol. 5, no. 11, p. 295, Nov. 9, 2021, ISSN: 2504-477X. DOI: 10.3390/jcs5110295 Accessed: Oct. 16, 2025. [Online]. Available: https://www.mdpi.com/2504-477X/5/11/295

[154] H. Pan, D. Yang, W. Qu, J. Li, and Y. Ke, "Process-dependent wrinkle formation for steered tow during automated fiber placement: Modeling and experimental verification," *Thin-Walled Structures*, vol. 180, p. 109 928, Nov. 2022, ISSN: 02638231. DOI: 10.1016/j.tws.2022.109928 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263823122005559

[155] A. Forghani et al., "Simulating Prepreg Tack in AFP Process," in *American Society for Composites 2017*, DEStech Publications, Inc., Nov. 15, 2017, ISBN: 978-1-60595-418-9. DOI: 10.12783/asc2017/15214 Accessed: Oct. 16, 2025. [Online]. Available: http://dpi-proceedings.com/index.php/asc32/article/view/15214

[156] A. Forghani et al., "Experimental Calibration of a Numerical Model of Prepreg Tack for Predicting AFP Process Related Defects," presented at the SAMPE, Diamond Bar, CA, United States, May 21, 2018. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002539

[157] R. Wehbe, B. Tatting, S. Rajan, R. Harik, M. Sutton, and Z. Gürdal, "Geometrical modeling of tow wrinkles in automated fiber placement," *Composite Structures*, vol. 246, p. 112 394, Aug. 2020, ISSN: 02638223. DOI: 10.1016/j.compstruct.2020.

112394 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822320302324

[158] M. Belhaj and M. Hojjati, "Wrinkle formation during steering in automated fiber placement: Modeling and experimental verification," *Journal of Reinforced Plastics and Composites*, vol. 37, no. 6, pp. 396–409, Mar. 2018, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/0731684417752872 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0731684417752872

[159] X. Li, S. R. Hallett, and M. R. Wisnom, "Modelling the effect of gaps and overlaps in automated fibre placement (AFP)-manufactured laminates," *Science and Engineering of Composite Materials*, vol. 22, no. 2, pp. 115–129, Mar. 1, 2015, ISSN: 2191-0359, 0792-1233. DOI: 10.1515/secm-2013-0322 Accessed: Oct. 16, 2025. [Online]. Available: https://www.degruyter.com/document/doi/10.1515/secm-2013-0322/html

[160] M. H. Nguyen, A. A. Vijayachandran, P. Davidson, D. Call, D. Lee, and A. M. Waas, "Effect of automated fiber placement (AFP) manufacturing signature on mechanical performance of composite structures," *Composite Structures*, vol. 228, p. 111 335, Nov. 2019, ISSN: 02638223. DOI: 10.1016/j.compstruct.2019.111335 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822319317672

[161] O. Falcó, J. Mayugo, C. Lopes, N. Gascons, and J. Costa, "Variable-stiffness composite panels: Defect tolerance under in-plane tensile loading," *Composites Part A: Applied Science and Manufacturing*, vol. 63, pp. 21–31, Aug. 2014, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2014.03.022 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X14001018

[162] R. Anay et al., "An experimental investigation concerning the effect of AFP defects on progressive damage in CFRP coupons," *Composite Structures*, vol. 279, p. 114 725, Jan. 2022, ISSN: 02638223. DOI: 10.1016/j.compstruct.2021.114725 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822321011776

[163] D. Del Rossi, V. Cadran, P. Thakur, M. Palardy-Sim, M. Lapalme, and L. Lessard, "Experimental investigation of the effect of half gap/half overlap defects on the strength of composite structures fabricated using automated fibre placement (AFP)," *Composites Part A: Applied Science and Manufacturing*, vol. 150, p. 106 610, Nov. 2021, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2021.106610 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X21003286

[164] K. Fayazbakhsh, M. Arian Nik, D. Pasini, and L. Lessard, "Defect layer method to capture effect of gaps and overlaps in variable stiffness laminates made by Automated Fiber Placement," *Composite Structures*, vol. 97, pp. 245–251, Mar. 2013, ISSN:

02638223. doi: 10.1016/j.compstruct.2012.10.031 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822312005272

[165] M. Ghayour, M. Hojjati, and R. Ganesan, "Induced defect layer method to characterize the effect of fiber tow gaps for the laminates manufactured by automated fiber placement technique," *Journal of Composite Materials*, vol. 55, no. 27, pp. 4011–4027, Nov. 2021, issn: 0021-9983, 1530-793X. doi: 10.1177/00219983211031649 Accessed: Oct. 19, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/00219983211031649

[166] M. Yi, F. Liu, W. Zhang, and X. Ding, "Formation Mechanism and Modeling Method of Wrinkling Defects in Variable Angle Tow Steering Fiber Placement," *Aerospace*, vol. 9, no. 10, p. 620, Oct. 19, 2022, issn: 2226-4310. doi: 10.3390/aerospace9100620 Accessed: Oct. 19, 2025. [Online]. Available: https://www.mdpi.com/2226-4310/9/10/620

[167] T. Tay, G. Liu, V. Tan, X. Sun, and D. Pham, "Progressive Failure Analysis of Composites," *Journal of Composite Materials*, vol. 42, no. 18, pp. 1921–1966, Sep. 2008, issn: 0021-9983, 1530-793X. doi: 10.1177/0021998308093912 Accessed: Oct. 19, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0021998308093912

[168] A. Orifici, I. Herszberg, and R. Thomson, "Review of methodologies for composite material modelling incorporating failure," *Composite Structures*, vol. 86, no. 1–3, pp. 194–210, Nov. 2008, issn: 02638223. doi: 10.1016/j.compstruct.2008.03.007 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822308000810

[169] A. De Luca, F. Caputo, and Department of Industrial and Information Engineering, University of Campania L. Vanvitelli, Via Roma 29, 81031 Aversa, Italy., "A review on analytical failure criteria for composite materials," *AIMS Materials Science*, vol. 4, no. 5, pp. 1165–1185, 2017, issn: 2372-0484. doi: 10.3934/matersci.2017.5.1165 Accessed: Oct. 19, 2025. [Online]. Available: http://www.aimspress.com/article/10.3934/matersci.2017.5.1165

[170] W. Zhou, Q. Cheng, Q. Xu, W. Zhu, and Y. Ke, "Deformation and fracture mechanisms of automated fiber placement pre-preg laminates under out-of-plane tensile loading," *Composite Structures*, vol. 255, p. 112 948, Jan. 2021, issn: 02638223. doi: 10.1016/j.compstruct.2020.112948 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822320328749

[171] M. Ghayour, R. Ganesan, and M. Hojjati, "Flexural response of composite beams made by Automated Fiber Placement process: Effect of fiber tow gaps," *Composites Part B: Engineering*, vol. 201, p. 108 368, Nov. 2020, issn: 13598368. doi: 10.1016/j.compositesb.2020.108368 Accessed: Oct. 19, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S135983682033417X

[172] S. Nikbakt, S. Kamarian, and M. Shakeri, "A review on optimization of composite structures Part I: Laminated composites," *Composite Structures*, vol. 195, pp. 158–185, Jul. 2018, ISSN: 02638223. DOI: 10.1016/j.compstruct.2018.03.063 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822317336279

[173] H. Ghiasi, K. Fayazbakhsh, D. Pasini, and L. Lessard, "Optimum stacking sequence design of composite materials Part II: Variable stiffness design," *Composite Structures*, vol. 93, no. 1, pp. 1–13, Dec. 2010, ISSN: 02638223. DOI: 10.1016/j.compstruct.2010.06.001 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822310001947

[174] A. Sabido, L. Bahamonde, R. Harik, and M. J. Van Tooren, "Maturity assessment of the laminate variable stiffness design process," *Composite Structures*, vol. 160, pp. 804–812, Jan. 2017, ISSN: 02638223. DOI: 10.1016/j.compstruct.2016.10.081 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822316314428

[175] G. G. Lozano, A. Tiwari, C. Turner, and S. Astwood, "A review on design for manufacture of variable stiffness composite laminates," *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 230, no. 6, pp. 981–992, Jun. 2016, ISSN: 0954-4054, 2041-2975. DOI: 10.1177/0954405415600012 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0954405415600012

[176] D. Punera and P. Mukherjee, "Recent developments in manufacturing, mechanics, and design optimization of variable stiffness composites," *Journal of Reinforced Plastics and Composites*, vol. 41, no. 23–24, pp. 917–945, Dec. 2022, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/07316844221082999 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/07316844221082999

[177] R. Harik, J. HALBRITTER, and A. BRASINGTON, "Computer-aided process planning (capp) for automated fiber placement (afp) manufacturing," U.S. Patent 20240051242A1, Feb. 15, 2024. Accessed: Oct. 16, 2025. [Online]. Available: https://patents.google.com/patent/US20240051242A1/en

[178] A. Brasington, J. Halbritter, R. Wehbe, and R. Harik, "Bayesian optimization for process planning selections in automated fiber placement," *Journal of Composite Materials*, vol. 56, no. 28, pp. 4275–4296, Dec. 2022, ISSN: 0021-9983, 1530-793X. DOI: 10.1177/00219983221129010 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/00219983221129010

[179] W. Qu et al., "Automated fiber placement path generation method based on prospective analysis of path performance under multiple constraints," *Composite Structures*, vol. 255, p. 112 940, Jan. 2021, ISSN: 02638223. DOI: 10.1016/j.compstruct.2020.

112940 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S026382232032866X

[180] W. Qu et al., "Placement suitability analysis of automated fiber placement on curved surfaces considering the influence of prepreg tow, roller and AFP machine," *Composite Structures*, vol. 262, p. 113 608, Apr. 2021, ISSN: 02638223. DOI: 10.1016/j.compstruct.2021.113608 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0263822321000696

[181] P. Debout, H. Chanal, and E. Duc, "Tool path smoothing of a redundant machine: Application to Automated Fiber Placement," *Computer-Aided Design*, vol. 43, no. 2, pp. 122–132, Feb. 2011, ISSN: 00104485. DOI: 10.1016/j.cad.2010.09.011 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0010448510001788

[182] N. Swingle, A. Brasington, J. Halbritter, and R. Harik, "Automated Fiber Placement Laminate Level Optimization for Mitigation of Through Thickness Defect Stacking," *SAMPE JOURNAL*, vol. 59, no. 6, pp. 32–47, Nov. 2023, ISSN: 00911062, 00911062. DOI: 10.33599/SJ.v59no6.03 Accessed: Oct. 17, 2025. [Online]. Available: https://www.nasampe.org/store/viewproduct.aspx?ID=23120106

[183] N. C. Swingle, A. R. Brasington, J. A. Halbritter, and R. Harik, "Implementation of Combinatorial Optimization Techniques for Automated Fiber Placement Through Thickness Defect Stack-Up Minimization," presented at the SAMPE Conference and Exhibition, Seattle, WA. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20230000366

[184] T. Aized and B. Shirinzadeh, "Robotic fiber placement process analysis and optimization using response surface method," *The International Journal of Advanced Manufacturing Technology*, vol. 55, no. 1–4, pp. 393–404, Jul. 2011, ISSN: 0268-3768, 1433-3015. DOI: 10.1007/s00170-010-3028-1 Accessed: Oct. 16, 2025. [Online]. Available: http://link.springer.com/10.1007/s00170-010-3028-1

[185] T. Yu, Y. Shi, X. He, C. Kang, B. Deng, and S. Song, "Optimization of Parameter Ranges for Composite Tape Winding Process Based on Sensitivity Analysis," *Applied Composite Materials*, vol. 24, no. 4, pp. 821–836, Aug. 2017, ISSN: 0929-189X, 1573-4897. DOI: 10.1007/s10443-016-9553-9 Accessed: Oct. 16, 2025. [Online]. Available: http://link.springer.com/10.1007/s10443-016-9553-9

[186] Z. Han, S. Sun, Z. Shao, and H. Fu, "Multiscale Collaborative Optimization of Processing Parameters for Carbon Fiber/Epoxy Laminates Fabricated by High-Speed Automated Fiber Placement," *Advances in Materials Science and Engineering*, vol. 2016, pp. 1–14, 2016, ISSN: 1687-8434, 1687-8442. DOI: 10.1155/2016/5480352 Accessed: Oct. 16, 2025. [Online]. Available: https://www.hindawi.com/journals/amse/2016/5480352/

[187] S. Sun, Z. Han, J. Zhang, H. Jin, and Y. Wang, "Multiscale collaborative process optimization method for automated fiber placement," *Composite Structures*, vol. 259, p. 113 215, Mar. 2021, ISSN: 02638223. DOI: 10.1016/j.compstruct.2020.113215 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S026382232033141X

[188] Z. Han, S. Sun, H. Fu, and Y. Fu, "Multi-Scale Low-Entropy Method for Optimizing the Processing Parameters during Automated Fiber Placement," *Materials*, vol. 10, no. 9, p. 1024, Sep. 3, 2017, ISSN: 1996-1944. DOI: 10.3390/ma10091024 Accessed: Oct. 16, 2025. [Online]. Available: https://www.mdpi.com/1996-1944/10/9/1024

[189] F. O. Sonmez and M. Akbulut, "Process optimization of tape placement for thermoplastic composites," *Composites Part A: Applied Science and Manufacturing*, vol. 38, no. 9, pp. 2013–2023, Sep. 2007, ISSN: 1359835X. DOI: 10.1016/j.compositesa.2007.05.003 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X07000814

[190] J. Jiang, Y. He, and Y. Ke, "Pressure distribution for automated fiber placement and design optimization of compaction rollers," *Journal of Reinforced Plastics and Composites*, vol. 38, no. 18, pp. 860–870, Sep. 2019, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/0731684419850896 Accessed: Oct. 16, 2025. [Online]. Available: https://journals.sagepub.com/doi/10.1177/0731684419850896

[191] Y. Li et al., "Numerical and experimental study of the viscosity and mechanical parameters effect on tow feeding quality during automated fiber placement," *Journal of Manufacturing Processes*, vol. 96, pp. 31–40, Jun. 2023, ISSN: 15266125. DOI: 10.1016/j.jmapro.2023.04.048 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1526612523004012

[192] D. Maurer and P. Mitschang, "Laser-powered tape placement process – simulation and optimization," *Advanced Manufacturing: Polymer & Composites Science*, vol. 1, no. 3, pp. 129–137, Jul. 3, 2015, ISSN: 2055-0340, 2055-0359. DOI: 10.1080/20550340.2015.1114798 Accessed: Oct. 16, 2025. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/20550340.2015.1114798

[193] Y. Li et al., "Modeling of tow tension fluctuations and parameter optimization during the stable transfer phase for automated fiber placement," *Tribology International*, vol. 183, p. 108 429, May 2023, ISSN: 0301679X. DOI: 10.1016/j.triboint.2023.108429 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0301679X23002165

[194] D. Heider, M. J. Piovoso, and J. W. Gillespie, "A neural network model-based open-loop optimization for the automated thermoplastic composite tow-placement system," *Composites Part A: Applied Science and Manufacturing*, vol. 34, no. 8, pp. 791–799, Aug. 2003, ISSN: 1359835X. DOI: 10.1016/S1359-835X(03)00120-9

Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359835X03001209

[195] A. Brasington, J. Halbritter, M. Godbold, M. Kirkpatrick, C. Sacco, and R. Harik, "Mapping of multimodality data for manufacturing analyses in automated fiber placement," *Composites Part B: Engineering*, vol. 263, p. 110 881, Aug. 2023, ISSN: 13598368. DOI: 10.1016/j.compositesb.2023.110881 Accessed: Oct. 16, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1359836823003840

[196] A. Noevere and C. Collier, "Mapping Manufacturing Data for Stress Analysis of Automated Fiber Placement Structures," in *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 8, 2018, ISBN: 978-1-62410-532-6. DOI: 10.2514/6.2018-0228.c1 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2018-0228.c1

[197] A. Noevere, C. Collier, and R. Harik, "Integrated Design and Manufacturing Analysis for Automated Fiber Placement Structures," presented at the SAMPE, Diamond Bar, CA, United States, May 20, 2019. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002674

[198] A. T. Noevere, C. S. Collier, R. Harik, and J. Halbritter, "Development of a Design for Manufacturing Tool for Automated Fiber Placement Structures," presented at the American Institute of Aeronautics and Astronautics (AIAA), Reston, VA, United States, Jan. 7, 2019. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002442

[199] A. T. Noevere and C. Collier, "Design for Manufacturing Tool for Automated Fiber Placement Structures – Verification and Validation," in *AIAA Scitech 2020 Forum*, Orlando, FL: American Institute of Aeronautics and Astronautics, Jan. 6, 2020, ISBN: 978-1-62410-595-1. DOI: 10.2514/6.2020-1477 Accessed: Oct. 16, 2025. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2020-1477

[200] A. Brasington, J. Halbritter, A. Noevere, and R. Harik, "Integration of Design Data Into Automated Fiber Placement Process Planning Metrics," presented at the Society of Advanced Materials and Process Engineering (SAMPE), Seattle, WA, 2023. Accessed: Oct. 16, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20230000352

[201] J. Moore, A. Colvin, S. Ghose, and B. Johnson, "Design for Manufacturing: Laminate Focused Design and Analysis Tools for Automated Composites Manufacturing," in *SAMPE 2019 - Charlotte, NC*, SAMPE, Apr. 11, 2019. DOI: 10.33599/nasampe/s.19.1478 Accessed: Oct. 16, 2025. [Online]. Available: https://sampe.knack.com/technical-papers#home/technicalpaperdetails/5cb7822a71efac0cb6e45a1c/

[202] B. Hasenjaeger, "Programming and Simulating Automated Fiber Placement (AFP) CNC Machines," *SAMPE Journal*, vol. 49, no. 6, pp. 7–12, Nov. 2013. [Online].

Available: https://www.electroimpact.com/whitepapers/programming_and_simulating_automated_fiber_placement_cnc_machines.pdf

[203] J. Halbritter and R. Harik, “Computer Aided Process Planning for Automated Fiber Placement,” in *SAMPE neXus Proceedings*, Jan. 7, 2021.

[204] R. Saaty, “The analytic hierarchy process—what it is and how it is used,” *Mathematical Modelling*, vol. 9, no. 3–5, pp. 161–176, 1987, ISSN: 02700255. DOI: 10.1016/0270-0255(87)90473-8 Accessed: Feb. 4, 2026. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/0270025587904738

[205] C. Collier and S. Jones, “Unified Analysis of Aerospace Structures through Implementation of Rapid Tools into a Stress Framework,” in *AIAA Scitech 2020 Forum*, Orlando, FL: American Institute of Aeronautics and Astronautics, Jan. 6, 2020, ISBN: 978-1-62410-595-1. DOI: 10.2514/6.2020-1478 Accessed: Feb. 5, 2026. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2020-1478

[206] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient Global Optimization of Expensive Black-Box Functions,” *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, Dec. 1998, ISSN: 0925-5001, 1573-2916. DOI: 10.1023/A:1008306431147 Accessed: Feb. 10, 2026. [Online]. Available: https://link.springer.com/10.1023/A:1008306431147

[207] R. Brent, *Algorithms for Minimization without Derivatives* (Dover Books on Mathematics). Dover Publications, 2013, ISBN: 978-0-486-14368-2. [Online]. Available: https://books.google.com/books?id=AITCAgAAQBAJ

[208] W. Press, *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Numerical Recipes: The Art of Scientific Computing). Cambridge University Press, 2007, ISBN: 978-0-521-88068-8. [Online]. Available: https://books.google.com/books?id=1aAOdzK3FegC

[209] B. Böckl, A. Wedel, A. Misik, and K. Drechsler, “Effects of defects in automated fiber placement laminates and its correlation to automated optical inspection results,” *Journal of Reinforced Plastics and Composites*, vol. 42, no. 1–2, pp. 3–16, Jan. 2023, ISSN: 0731-6844, 1530-7964. DOI: 10.1177/07316844221093273 Accessed: Mar. 22, 2026. [Online]. Available: https://journals.sagepub.com/doi/10.1177/07316844221093273

[210] I. Newton, *Philosophiæ Naturalis Principia Mathematica*. 1848. Accessed: Apr. 5, 2026. [Online]. Available: https://archive.org/details/newtonsprincipia00newt_0

[211] D. Bertsekas, *Nonlinear Programming*. Athena Scientific, Sep. 1, 2016, 879 pp., ISBN: 978-1-886529-05-2. Google Books: rC1EEAAAQBAJ.

[212] A. Yousefpour et al., “Enabling Responsive Real-Time Inspection of the Automated Fiber Placement Process,” in *SAMPE 2020 | Virtual Series*, NA SAMPE, 2020. DOI:

10.33599/nasampe/s.20.0258 Accessed: Jan. 27, 2026. [Online]. Available: https://www.digitallibrarynasampe.org/data/pdfs/s2020_pdfs/TP20-0000000258.pdf

[213] J. Halbritter, "Leveraging Automated Fiber Placement Computer Aided Process Planning Framework for Defect Validation and Dynamic Layup Strategies," Ph.D. dissertation, University of South Carolina, Apr. 1, 2023. [Online]. Available: https://scholarcommons.sc.edu/etd/7212