

Uniform Coverage of Simple Surfaces Embedded in \mathbb{R}^3 for Auto-Body Painting

Prasad N. Atkar, David C. Conner, Aaron Greenfield, Howie Choset, and
Alfred A. Rizzi

Carnegie Mellon University, Pittsburgh PA 15213, USA

Abstract. We develop a procedure for automated trajectory generation for robotic spray painting applications. Painting requires that the spray gun deposit paint at each point on the target surface such that the variation of the resultant paint deposition is within acceptable limits; we term this the uniform coverage problem. To understand the key issues in the uniform coverage problem, we consider surface patches that are geodesically convex and topologically simple to represent realistic automotive surfaces. Our goal is to understand the relationship between the spray gun trajectory and the following output characteristics: deposition uniformity, cycle time, and paint waste. Our planning approach decomposes the coverage trajectory generation problem into three subproblems: selection of the start curve, selection of the speed profiles along each pass, and selection of the spacing between the passes. Using concepts such as area magnification, the Gauss-Bonnet theorem from differential geometry, and standard optimization procedures, we present procedures to solve each subproblem independently of the others. Finally, we demonstrate our trajectory planning procedures in simulation as well as experimentally on simple surfaces that approximate real automobile surfaces.

1 Introduction

Today, robots are widely used for spray-painting in the automotive industry. Among the different kinds of spray painting mechanisms available, electrostatic rotating bell (ESRB) atomizers are one of the most popular. However, the distribution of paint generated by an ESRB atomizer is relatively complex. This complexity is further compounded as the spray gun moves over non-planar automotive surfaces, making the task of planning paths for such atomizers challenging. Paint specialists typically produce the coverage paths based on their experience, often requiring 3 to 5 months to completely plan trajectories on a new automobile model. This programming time is a critical bottleneck in the “concept-to-consumer” timeline for bringing a new automobile to the market. Automating the process of path planning will help the paint specialists reduce this programming time significantly by offering them reasonable guidelines for effective paths.

In this work, we develop procedures for automated generation of spray gun trajectories to optimize the output characteristics over simple surfaces. By “simple surfaces,” we mean surfaces that are geodesically convex and have no holes. A surface is termed geodesically convex if the shortest curve joining any two points on the surface is a geodesic (see Fig. 1(a) and 1(b)). We use such surfaces to closely approximate the auto-body surface parts.




Fig. 1. (a) A geodesically convex bent sheet; here the shortest curve joining any two points is a geodesic. (b) the surface is not geodesically convex because the shortest curve in the surface joining two points A and B is not a geodesic. (c) We generate a coverage path on arbitrary simple surfaces by selecting a start curve and offsetting it sideways within the surface to generate new passes.

To generate a spray gun trajectory termed a *coverage trajectory* that uniformly covers (i.e., deposits paint on) the target surface, the path planning algorithm must determine the orientation of the passes in the path, the spacing between them and the speed along the passes. The dimensionality of the variables required to specify these three entities is huge. Therefore, global optimization procedures that attempt to determine all coverage variables simultaneously are computationally expensive and often not practical.

We seek to make the optimization of coverage variables tractable by decomposing the coverage problem into three relatively independent subproblems: i) selection of a “seed” pass termed the *start curve* on the surface (see Fig. 1(c)), ii) selection of the speed profile along a given pass, and iii) selection of the optimum spacing between a given pass and its adjacent pass. Based on these three subproblems, we generate the coverage trajectory using the procedure described in Algorithm 1.

Data : Target surface CAD model, parameters for deposition model
Result : Spray gun trajectory
 Select a pass termed the start curve on the surface;
 Optimize end-effector speed along the start curve;
repeat
 | Offset the most recently generated pass within the surface to its “right”
 | side to obtain a new offset pass;
 | Optimize end-effector speed along the new offset pass;
until *Offset pass lies completely outside the surface*;
 Designate the start curve as the most recently generated pass;
repeat
 | Offset the most recently generated pass within the surface to its “left” side
 | to obtain a new offset pass;
 | Optimize end-effector speed along the new offset pass;
until *Offset pass lies completely outside the surface*;

Algorithm 1: Path Planning algorithm

The rationale behind decomposing the coverage problem into the three subproblems is that each subproblem can be solved in a reasonable amount of

time independently from the others. The cycle time for the coverage trajectory is primarily dependent on the orientation of the passes in the path. The selection of a particular start curve essentially determines the orientation of the rest of the passes, and consequently the cycle time required. The structure of the paint profiles measured in the direction orthogonal to a pass is relatively consistent along the direction of the pass (see Fig. 2(a)). Therefore, the uniformity of paint deposition can be seen as having two components: 1) uniformity along the direction of the passes, and 2) uniformity in the direction orthogonal to the passes. Speed optimization attempts to produce similar paint profiles along a pass, whereas the optimization of the spacing between passes, termed the *index width*, attempts to overlap the paint profiles of two adjacent passes appropriately to produce uniform paint deposition in the direction orthogonal to the passes.

In this work, we continue our prior work on deposition modelling and start curve selection, briefly discussed in Sections 2.2 and 2.3 respectively, by providing procedures for speed optimization and index width optimization. We formulate the speed optimization problem in Section 3. We address the index optimization problem in Section 4 for surfaces with increasing geometric complexity: planar surfaces, extruded sheets and surfaces with non-zero Gaussian curvature. Finally, we demonstrate our coverage procedures for a variety of surfaces in simulation as well as experimentally in Section 5.

2 Prior Work

2.1 Related Work

Most prior researchers typically focus on a particular subproblem for the path generation procedure: 1) start curve selection, 2) speed optimization, or 3) index width optimization. However, only a few researchers study all the three problems together.

For start curve selection, most prior researchers [1–4] select the pass orientation that aligns with one of the faces of a bounding box that fits the surface, while choosing the relative position of the start curve arbitrarily. Such an approach of start curve selection implicitly tries to minimize cycle time, but does not consider the effects of the relative position of the start curve on paint uniformity and can lead to poor uniformity results. For a planar surface, Huang [5] gives a sophisticated approach for minimizing the cycle time for coverage by minimizing number of passes in the path. Kim and Sarma [6] use vector fields to choose the pass orientation that minimizes the cycle time with an implicit constraint on the paint deposition uniformity. The effect of the vector field orientation on the paint uniformity is not explicitly considered. In an attempt to maximize uniformity, Smith *et al.* [7] select the orientation of the passes by determining the section plane that is maximally orthogonal to the surface. Their approach does not consider minimizing the cycle time or the paint waste and is sensitive to small curvature changes in the target surface.

One of the most general frameworks for efficiently optimizing the speed profile of a coverage path is given by Antonio and Ramabhadran [8]. Their work assumes that the coverage path is already known and the deposition




Fig. 2. (a) A typical deposition model and the paint profile across the paint swath on a flat planar sheet. (b) Our current model determines the deposition on any given point on an arbitrary surface using concept of area magnification under the linear projection model.

model is either a bivariate Cauchy or a Gaussian distribution applied to a flat panel. Kim and Sarma [6] also present a speed optimization model for a coverage path generation framework based on vector fields with a particular focus on the process cycle time.

Most prior path planning approaches consider simplistic deposition models such as circular [1], parabolic [4,9,10], or beta distributions [2]. Likewise, most of these approaches make first order approximations to the surface geometry, thus limiting their use for realistic auto surfaces. In these cases, the selection of index widths between adjacent passes is easier [1–4,6,7], but fails to capture the realistic scenario.

Finally, commercially available path planning systems such as RobCADTM¹ project a user-specified planar path on to the target surface and yield the output characteristics of the resultant coverage path. The limited scope for the specification of the deposition model and the requirement for manual specification of the coverage path limit the utility of such software tools.

2.2 Our Prior Work: Deposition Modelling

In automated path planning systems, it is necessary to have a paint deposition model that can predict paint deposition on an arbitrary surface with a reasonable accuracy to effectively determine the suitability of a given coverage trajectory. In [11], we develop a simple model that provides a significant improvement in paint prediction over the earlier models while retaining sufficient tractability for use in our planning tools. The deposition model captures the shape of the paint distribution from a spray gun in an analytical representation. The deposition model is composed of one bivariate and two revolved Gaussians [11,12]. We extract the parameters for the deposition model by applying data fitting techniques to the experimental data obtained by painting flat panels (see Fig. 2(a)). We then determine the deposition at any point on a given arbitrary curved surface by using the area magnification concept from differential geometry (see Fig. 2(b)).

¹ A product of Tecnomatix Technologies Ltd.

For simplicity, our current model assumes that the paint particles flow in a straight line after leaving the spray gun nozzle. While this is an incorrect assumption, our experience shows us that this assumption is reasonable for surfaces with low curvature. In the future, we plan to extend our techniques to include more realistic projection models. We would like to emphasize that although deposition models are required by the planning procedure, the coverage planning procedures we develop in this work are independent of any particular deposition model.

2.3 Our Prior Work: Start Curve Selection

The choice of start curve impacts the entire coverage path and its cycle time. Unfortunately, the search-space for the start curve selection problem is huge. In [13], we select a start curve such that the resultant coverage path attempts to minimize the cycle time and yields the desired paint deposition uniformity. Uniformity of paint deposition suffers severely in the worst possible case where the passes in the coverage path self-intersect. The more the start curve bends “sideways” in the surface, or equivalently, the higher the geodesic curvature of the start curve, the greater the risk that the subsequent offset curves will exhibit a self-intersection [13,14]. Therefore, to minimize the possibility of self-intersections on the offset curves, we seek to select start curves that have low geodesic curvature, or ideally, are themselves geodesic. This restriction on the choice of start curve to the family of geodesics reduces the start curve selection problem to determining two variables: 1) the orientation of the (geodesic) start curve, and 2) the relative position of start curve with respect to the surface boundary.

Our path construction procedure (see Algorithm 1), first offsets the start curve and then continues to offset newly generated passes until the surface is covered completely. To ensure that the subsequent offset passes are also free from self-intersections, we must ensure that the geodesic curvature on any newly generated pass is also minimal. Thus, ideally we want to select a start curve such that all the resultant passes are geodesics. However, on a surface with non-zero Gaussian curvature, the offset of a geodesic curve is, in general, not a geodesic. In such a case, the position of the start curve relative to the boundary dictates the resultant geodesic curvature on the offset curves.

To examine the effect of surface Gaussian curvature K on the geodesic curvature κ_g of the offset curves, we apply the Gauss-Bonnet theorem to the region B bounded between the start curve, C_{st} , and its offset curve, C_{of} , (see Fig. 3(a)) and arrive at

$$\int_{C_{of}} \kappa_g = \int_B K + \int_{C_{st}} \kappa_g. \quad (1)$$

Equation 1 tells us that the more the surface bounded between the offset curve and the geodesic start curve bends, the more the geodesic curvature of the offset curve increases. On a surface where the sign of the Gaussian curvature stays the same, the further away we place the offset curve from the start curve, the higher the geodesic curvature of the offset curve is. Therefore, the possibility of self-intersection is maximal on the bounding passes




Fig. 3. (a) Application of Gauss-Bonnet theorem to the region bounded between the start curve and the offset curve. (b) We select the orientation of the start curve that minimizes the surface altitude in the indexing direction, and consequently the cycle time.

(see Fig. 3(a)). To minimize the geodesic curvature on the bounding passes, Equation 1 tells us that the start curve should divide the surface into two parts such that the integral of the Gaussian curvature is equal over each part. This approach determines the relative position of the start curve on the surface as defined by the *geodesic Gaussian curvature divider* curve.

To determine the orientation of the start curve, we approximate the geodesic start curve by a curve of planar intersection that is also a Gaussian curvature divider. To ensure that such an approximation is sufficiently close to being a geodesic, we require that the normal to the section plane chosen for intersection is orthogonal to the average target surface normal. We generalize the concept of surface “altitude” from planar surfaces [5] to non-planar surfaces, and determine the orientation of the section plane that minimizes the surface altitude in the direction normal to the section plane (see Fig. 3(b)). The resultant orientation of the start curve yields a coverage path that minimizes the number of turns in the coverage path (for the given constraint on section plane normal) and equivalently, the cycle time and the paint waste.

3 Speed Optimization

In the automobile industry, constant speed trajectories are typically used. Because of the relatively large size of the ESRB deposition pattern with respect to a typical automotive surface, constant speed profiles typically require long oversprays to ensure that the boundary effects, which produce non-uniform paint deposition near surface boundaries, are minimized. If there is a restriction on the maximum amount of paint waste, shorter oversprays could be used, but not without a compromise in the paint deposition uniformity. Additionally, on non-planar surfaces, the changes in the surface curvature along the pass result in non-uniform paint deposition along the pass. Speed optimization attempts to compensate for these curvature related results and improves the uniformity of paint deposition in the direction of the passes.

Unlike prior approaches [6,8] that use speed optimization techniques over the entire path, our approach uses a semi-global method that optimizes speed profiles on a pass-by-pass basis. The underlying assumption in our approach




Fig. 4. Speed optimization attempts to minimize the variation of paint profiles in the direction of the pass.




Fig. 5. Graph of normalized standard deviation of paint deposition vs. the index width for a typical ESRB deposition pattern.

is that the speed optimization improves the uniformity in the direction of the passes and a local change in the speed profile affects the paint deposition uniformity on only a subset of the surface. Such an approach tries to combine the best of global and local optimization techniques.

Our implementation first samples a “band” of the surface around a given pass by constructing several “sub-offset” curves which are similar to the offset curves of the current pass, but with a much smaller spacing between them (see Fig. 4). Our objective here is to minimize the weighted sum of standard deviations of paint deposition along all sub-offset curves, thus minimizing the variations in the paint profile along the pass.

Our implementation first discretizes the given spray gun pass into n linear segments of length s_i . We assume that the speed of the spray gun tool center point (TCP) does not vary in a given segment i and is given by v_i . The speed profile along the pass is represented by the n -tuple $\mathbf{V} = \{v_i : i \in [1, n]\}$. The time t_i spent by the spray gun in each segment i of the pass is equal to $\frac{s_i}{v_i}$. Instead of optimizing the speed profile \mathbf{V} explicitly, we optimize the time profile $\mathbf{T} = \{t_i : i \in [1, n]\}$, and then calculate \mathbf{V} accordingly.

To determine the standard deviation of paint deposition along each sub-offset curve j , we discretize the j th sub-offset curve into m_j linear segments of length s_{m_j} . Let \mathbf{D}_j represent the matrix $[d_{e,f}]$ which gives the deposition flux on segment e of the j th sub-offset curve when the spray gun tool center point lies at the center of segment f of the current pass.

For a given time profile \mathbf{T} , the resultant paint deposition on the j th sub-offset curve is given by $\mathbf{D}_j\mathbf{T}$. Note that $\mathbf{D}_j\mathbf{T}$ is a m_j -tuple whose i th component gives the deposition on i th segment of j th sub-offset curve. Let S_j be the sum of all elements of vector $\mathbf{D}_j\mathbf{T}$. The average paint deposition k_j along the j th sub-offset curve is then equal to $k_j = \frac{S_j}{m_j}$. The normalized standard deviation of paint deposition along j th sub-offset curve is then computed as $\|\frac{\mathbf{D}_j\mathbf{T} - \mathbf{K}_j}{k_j}\|$, where \mathbf{K}_j is a m_j -tuple with elements of constant values equal to k_j .

The constraints on the maximum and minimum speed are specified at each segment along the spray gun pass by requiring $\frac{s_i}{v_{\max}} \leq t_i \leq \frac{s_i}{v_{\min}} \quad \forall i$.

The acceleration constraint has a non-linear form in each segment which we linearize by using a conservative linear constraint that always satisfies the intended acceleration constraint. Here, the bounds for the maximum acceleration and deceleration are assumed to be a_{max} and $-a_{max}$ respectively.

Let $w_j = \frac{k_j}{\sum_i k_i}$ be the weight of the normalized standard deviation along j th sub-offset curve in the objective function. Then, the complete speed optimization problem is formulated as:

$$\begin{aligned} \min_{\mathbf{T}} \sum_j \frac{w_j}{k_j} \|\mathbf{D}_j \mathbf{T} - \mathbf{K}_j\| & \quad (2) \\ \text{such that } \mathbf{D}_0 \mathbf{T} = \mathbf{K}_0 & \\ \frac{s_i}{v_{max}} \leq t_i \leq \frac{s_i}{v_{min}}, \text{ and } -a_{max} \frac{s_{i+1} s_i^2}{v_{max}^3} \leq s_{i+1} t_i - s_i t_{i+1} \leq a_{max} \frac{s_{i+1} s_i^2}{v_{max}^3}. & \end{aligned}$$

At the beginning of the optimization procedure, we specify an initial time trajectory \mathbf{T}^{beg} that corresponds to a constant speed profile with speed v_{nom} , where v_{nom} is chosen according to the desired average paint thickness. Each vector \mathbf{K}_j is then calculated using \mathbf{T}^{beg} , that is, $\mathbf{K}_j = \mathbf{D}_j \mathbf{T}^{\text{beg}}$. We then have all the initial conditions required for the optimization, and execute the optimization algorithm. Without loss of generality, we assume that the 0th sub-offset curve is along the spray gun pass. The equality constraint $\mathbf{D}_0 \mathbf{T} = \mathbf{K}_0$ enforces that all speed profiles yield approximately the same average paint deposition on the surface.

The objective function in our speed optimization formulation is a quadratic form in the variable \mathbf{T} with all constraints (i.e., equality and non-equality) that are linear in \mathbf{T} . This quadratic constrained optimization can be solved using standard optimization techniques with superlinear convergence. Note that a fine path resolution $\{s_i\}$ is necessary for accurate uniformity estimation, but results in long running times due to large number of inequality constraints. We address this issue by using a fine path resolution for evaluation of the objective function, but optimizing the speed profile over only a sub-sample of non-contiguous segments. The speed at the rest of segments in the path is determined by assuming a linear variation in speed between the segments for which we optimize the speed. This approach solves the optimization problem in a relatively short amount of time, and yet is acceptable as the industrial robots are typically ‘‘taught’’ with only a few points.

4 Index Width Optimization

The goal of speed optimization is to produce acceptable paint uniformity along the direction of a pass. Given speed optimized passes, the remaining problem is to decide how to place passes next to one another on the surface – that is, select *index widths* between the passes. The objective is to select index widths such that the paint profiles of adjacent passes overlap appropriately and produce acceptable uniformity orthogonal to the direction of the passes. At the same time, the index widths should be as wide as possible in order to reduce the number of passes in the coverage path, thereby reducing the cycle time and paint waste. In this section, we present procedures for index

width optimization on surfaces with increasing geometric complexity: planar surfaces, extruded surfaces, and surfaces with non-zero Gaussian curvature.

4.1 Determining Index Widths on Planar Surfaces

On a planar surface, geodesics are simply straight lines. Therefore the start curve, chosen as described in Section 2.3, will be a straight line. The offsets of the start curve are parallel straight lines and accordingly the resultant coverage path consists of a family of parallel lines.

To study the effect of index width on paint deposition uniformity, we plot the graph of the normalized standard deviation of paint deposition versus index width between passes (see Fig. 5) [12]. To evaluate the paint uniformity, we consider the interactions between the deposition profile curves of a sufficiently large number of passes spaced at constant index width.

From the graph, we observe that there is a “sweet spot” of index width that corresponds to a local minimum of the standard deviation (around 525 mm index width for the atomizer whose deposition pattern was considered in Fig. 5). Painting the target surface at this higher index width is desirable because higher index widths reduce the process cycle time. Unfortunately, the sensitivity of the standard deviation of paint deposition to the index width is high at the sweet spot; in other words, small changes in index width at the sweet spot produce high variations in paint deposition uniformity. In order to ensure that the paint deposition uniformity over the surface is not sensitive to small changes in index width, we typically do not use the sweet spot spacing between the passes.

Our implementation defines the index width search range, \mathcal{SR} , as a finite collection of potential index widths taken from the closed interval $[w_{min}, w_{max}]$ sampled at an appropriate resolution w_{res} . The minimum index width, w_{min} , is chosen according to the cycle time consideration, while the maximum index width, w_{max} , is chosen according to the standard deviation sensitivity. To reduce the computational costs, we choose w_{res} as large as possible, yet sufficiently small to meaningfully capture the variation of uniformity as a function of index width. Next, from the index width search range, we determine the set of feasible index widths that yield the normalized standard deviation of paint deposition below the user specified limit.

We then establish a cost function over the set of feasible index widths by assigning costs inversely proportional to the index width, thus penalizing the process cycle time. The feasible index width that minimizes the cost function is chosen as the optimum index width. By the design of the cost function, the optimum index width is the largest index width that yields an acceptable uniformity. Since the planar surface locally appears the same everywhere, the same optimum index width is chosen for all passes, assuming that there is sufficient overspray.

4.2 Determining Index Widths on Extruded Sheets

To lift the index width selection framework from planar surfaces to non-planar surfaces, we first consider a special class of target surfaces – extruded sheets. Extruded surfaces have zero Gaussian curvature, and are non-planar

in general. Many automobile surfaces, such as doors, are designed based on extruded surfaces, making this a useful class of surfaces.

Although different choices of start curve orientation are available on the extruded surface as described in Section 2.3, in this section we assume that the passes are along the zero curvature direction (see Fig. 6). Thus, the surface curvature orthogonal to each pass varies from pass to pass. In general, the paint profiles along any two passes on an extruded surface are different due to the variation in surface curvature. As a result, the optimal spacing between passes varies as surface curvature changes.

Our path planning approach generates new passes by offsetting known passes as described in Algorithm 1. At a given instance of time, we term the known pass we are offsetting the *current pass*. We then number all known passes relative to the current pass along the indexing curve that is orthogonal to the current pass. Let the known passes be denoted by $\{C_{-m}, \dots, C_0\}$, where C_0 is the current pass and C_{-m} is the m^{th} known pass from the current pass. Similarly, we denote n number of *future passes*, whose locations are unknown, by $\{C_1, \dots, C_n\}$. We term the first future pass, C_1 as the *candidate pass*.

Let the index width between pass C_i and pass C_{i+1} , measured along the indexing curve, be denoted by w_i . For m known passes, we define the ordered collection of known index widths as $\mathcal{K} = \{w_{-m+1}, w_{-m+2}, \dots, w_{-1}\}$. Similarly, the ordered collection of potential index widths for the n unknown future passes is defined as $\mathcal{U} = \{w_0, w_1, \dots, w_{n-1}\}$.

We are interested in generating the candidate pass, or equivalently determine the index width w_0 between the current pass and the candidate pass, such that the resultant paint deposition on the optimization profile is acceptably uniform. For the size of a typical deposition pattern of an ESRB atomizer, the paint deposition uniformity on the optimization profile is a function of the index widths of subsequently generated future passes, in addition to index widths of known passes.

We require that each potential index width belongs to index width search range \mathcal{SR} as defined in Section 4.1. The ordered collection $\mathcal{W} \equiv \{\mathcal{K}, \mathcal{U}\}$ denotes index widths between all passes, known and future (unknown). For

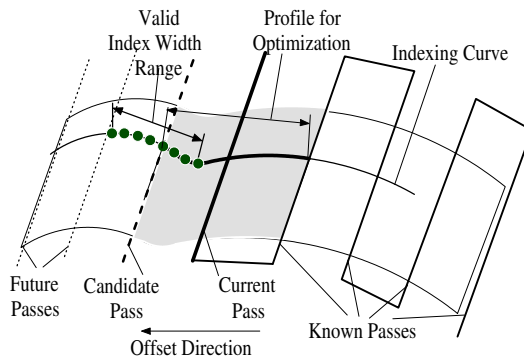


Fig. 6. Index width optimization on extruded sheets: the paint deposition uniformity is calculated along the indexing curve bounded between the previous pass and the candidate pass.

a given collection \mathcal{W} of index widths between all passes, we evaluate the normalized standard deviation, $u(\mathcal{W})$, of the paint deposition on the segment of the indexing curve bounded between pass \mathcal{C}_{-1} and the candidate pass \mathcal{C}_1 . We term this segment of the indexing curve bounded between passes \mathcal{C}_1 and \mathcal{C}_{-1} the *optimization profile*.

Our brute force approach varies the positions of the candidate and subsequently generated passes by using various collections of elements taken from \mathcal{SR} to compute \mathcal{U} . For each resultant collection of index widths \mathcal{W} , we evaluate the uniformity $u(\mathcal{W})$ over the optimization profile. We define the set of feasible candidate index widths, \mathcal{F} , as

$$\mathcal{F} = \{w'_0 : \exists \mathcal{W} = \{w_{-m+1}, \dots, w_0, \dots, w_{n-1}\}, w'_0 = w_0, u(\mathcal{W}) \leq u_{max}\}, (3)$$

where u_{max} is the maximum allowed normalized standard deviation. An index width in the feasible index set yields acceptable uniformity of paint deposition on the optimization profile for some set of future passes. In our implementation, for each value $w_0 \in \mathcal{SR}$, we vary the positions of future passes to find a set that yields acceptable uniformity on optimization profile. If such a set is found, we mark the index width w_0 as feasible. On the other hand, if the paint deposition uniformity over the optimization profile is not acceptable for any possible sets of future passes, the index width w_0 is marked as infeasible.

After the feasible index width set \mathcal{F} is determined, we establish a cost function for each index width in the feasible set, as in the planar case. Here, the cost function not only penalizes the smaller index widths, but also the difference between the current index width and the previous index width (i.e., $w_0 - w_{-1}$). This additional cost component ensures an appropriate balance between the deposition on the region between the current pass and the previous pass, and the deposition on the region between the candidate pass and the candidate pass. Then, the feasible index width that yields the minimal cost is chosen as the optimum index width, and accordingly the new candidate offset pass is generated.

4.3 Determining Index Widths on Surfaces with Non-zero Curvature

On surfaces with non-zero Gaussian curvature, the curvature of the surface, in general, changes not only as we move along an indexing curve (as in extruded surface case), but also as we move along a given pass. As such, the geometry of the indexing curve changes as we move along a given pass. Here, in order to determine the offset of the current pass, we sample the current pass at a finite number of “marker” points spaced at intervals based on the total curvature of the pass. At each marker point, we then determine the *indexing curve* as the intersection curve of a plane that is orthogonal to both the surface and the tangent to the pass at that marker point. We determine the optimum index width along each indexing curve and obtain the corresponding offset marker point by tracing along the indexing curve a distance equal to the optimum index width from the current marker point. The offset curve is then determined by interpolating between the collection of offset marker points.

To determine the optimum index width at each marker point, we first approximate the surface locally with a surface of extrusion generated by

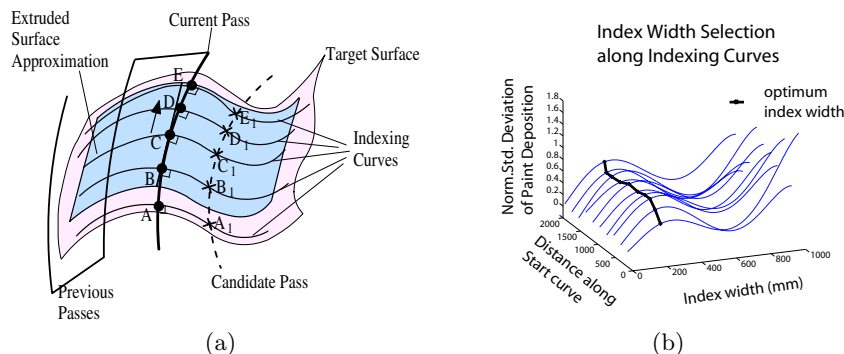


Fig. 7. Index width optimization on arbitrary surfaces: (a) at each of marker point A, B, C, D , and E , we approximate the surface with an extruded surface using the indexing curve, and determine the uniformity graph along the indexing curve. (b) The optimization procedure determines the optimum index width at each marker point. A_1, B_1, C_1, D_1 and E_1 are the corresponding offsets of the marker points.

extruding the indexing curve along the direction tangent to the pass at the marker point (see Fig. 7(a)). Employing the same approach used for the extruded sheets, we then determine the set of feasible index widths \mathcal{F}^i at each marker point i . For l marker points, we construct *index sets* formed by a combination operation by picking a single element from each \mathcal{F}^i at a time. That is, an index set is represented as $\mathcal{I} = \{w^1, w^2, \dots, w^l\}$, where $w^i \in \mathcal{F}^i$. Note that each index set, formed by choosing a different combination of elements from each \mathcal{F}^i , represents a different offset curve.

For each index set, we assign a cost that penalizes smaller index widths in the set, change in the index width from previous passes to current passes, as well as the variation of index width along the current pass. The last component of the cost function, in some sense, attempts to minimize the geodesic curvature of the offset curve.

Selecting the weights for the cost function components that produce satisfactory results for a general class of surfaces depends on the relationships between the costs for non-uniformity, cycle time and computational time; and sometimes this choice may not be readily apparent. In our implementation, we were principally motivated to reduce the occurrence of the zig-zag structure on the offset passes that not only renders the motion of the real-world end-effector impractical but may also lead to self-intersections of subsequent offset curves. An obvious way to achieve this goal is to assign heavy penalties for variation in the index width along the current pass. However, for simplicity in our implementation, we used a hard constraint that the index widths at all marker points should be equal. In other words, we consider only those offset passes which maintain a constant spacing from the current pass at all marker points.

This restriction of optimizing over only those index sets whose elements have equal values helps us to: 1) select the cost function with relative ease, 2) comply with the “constant index” offset assumption in selecting the start curve as discussed in Section 2.3, and 3) reduce the computational cost for index width selection for l marker points and index resolution w_{res} from

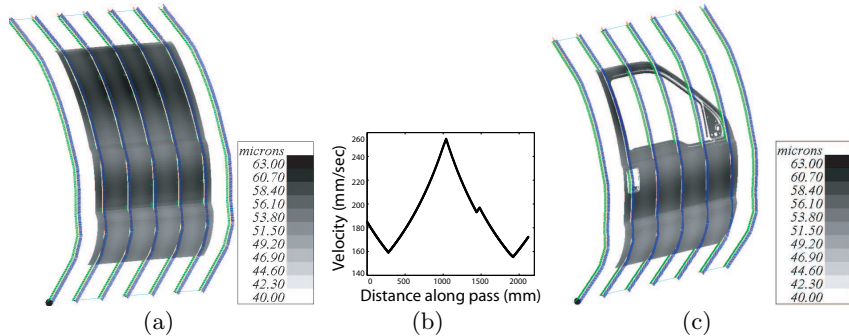


Fig. 8. (a) Simple surface approximation to the Ford Excursion door and the resultant paint deposition on the surface. The colormap shown is specifically designed to bring out the paint deposition variation. (b) A typical optimized speed profile along a spray gun pass. (c) Once the path on the approximation surface is generated, we simulate the paint deposition on the CAD model of a Ford Excursion door.

$O\left(\left(\frac{1}{w_{res}}\right)^{nl}\right)$ to $O\left(\left(l \frac{1}{w_{res}}\right)^n\right)$, where n (the number of future passes considered) is typically 4. We then select the optimum index set that minimizes the other two components of cost function from all possible constant-index sets (see Fig. 7(b)) and accordingly generate the candidate pass.

Note that although this “constant index” set solution that may not be globally optimal, it enables us to solve the index width selection problem satisfactorily within a practical amount of time. Moreover, the constant index set solution can serve as the starting point for a gradient-descent based method that considers all index sets for further improvement in uniformity.

5 Simulation and Experimental Results

To validate the utility of our coverage procedures, we generate spray gun trajectories using our planning tools on a variety of simple surfaces and study the effects on the resultant paint deposition uniformity. To evaluate paint deposition uniformity yielded by our generated trajectories, we simulate the paint deposition process on the corresponding surfaces. We also experimentally determine the paint uniformity on some real automobile surfaces. The surfaces we consider have varied geometric complexity including planar sheets, cylindrical surfaces, a door panel from a Ford Excursion and a fender from a Ford Crown Victoria. We model each of the two automobile surfaces by a slightly simplified single C^2 -continuous NURBS surface, by removing holes and merging multiple NURBS patches together from the corresponding CAD data. We then generate paths on the corresponding approximation surfaces, and use the same paths for experimentally painting the surfaces and for simulating paint deposition on the surface CAD models (see Fig. 8).

5.1 Speed Optimization Results

To examine the improvement in the resultant paint deposition uniformity using speed optimized trajectories over constant speed trajectories, we evaluate the resultant paint uniformity in each case over a few surfaces and list

the results in Table 1. We measure the uniformity in terms of normalized standard deviation of resultant paint deposition. The results show that in simulation, the speed optimization substantially improves the uniformity of resultant paint deposition. Also, the planar surface example shows that to produce similar uniformity, the overspray required by the speed optimized trajectories is shorter than that required by a constant speed trajectory.

Table 1. Speed optimization results

Surface	Simulation		Experiment	
	Const.Speed	Optimized	Const.Speed	Optimized
Excursion door	21.82 %	15.73 %	17.80 %	15.88 %
Approx. to Crown Victoria fender	16.62 %	11.92 %	-	-
Plane 100mm overspray	8.60 %	3.88 %	-	-
Plane 235mm overspray	3.88 %	-	-	-

The experimental data on the Ford Excursion door, however, shows that the speed optimized trajectory yields only a small improvement in the paint deposition uniformity over constant speed trajectory. Note that the speed optimization algorithm inherently relies on simulation of paint deposition to evaluate the objective function for a given speed profile. If there is a significant difference between simulated deposition and the experimental deposition, just as we observed, the effectiveness of speed optimization diminishes for the real-world application. More realistic deposition models will enable us to effectively use the speed optimization procedures in real systems.

5.2 Index Width Optimization Results

To study the effect of index width optimization on paint deposition uniformity, we consider a variety of target surfaces, generate passes for each of them and evaluate the paint deposition uniformity by simulating the paint deposition process (see Table 2). In each case, we set the maximum allowed normalized standard deviation at 4% and assume sufficient overspray to minimize the boundary effects. We observe that the index optimization yields the desired uniformity in cases where the surface curvature remains the same along the indexing curve. On the other hand, if the surface curvature changes along the indexing curve, the resultant paint deposition uniformity does not match the desired levels of uniformity. This is the case for Ford Excursion (horizontal passes) and the Ford Crown Victoria fender. The higher standard deviations observed on these surfaces are not unexpected because the index optimization process is, at best, semi-global in nature, and cannot globally guarantee the desired level of uniformity over the entire surface.

We also studied the effect of index optimization on paint uniformity experimentally by painting the middle and lower portion of a Ford Excursion door. During the experiments, our planning tools generated the index optimized path with a sideways overspray pass near the door bottom. In simulation, this path yielded a paint deposition uniformity of 7.16% on the door. However, because of robot workspace constraints, we had to manually remove the bottom overspray pass from the coverage path. For the coverage path

without the overspray pass, the uniformity of resultant paint deposition was 13.13% in simulation, whereas experimentally it was measured to be 16.5%. If we had optimized the index width with the constraint that there would be no bottom overspray pass, the resultant paint deposition uniformity could be improved to 9.01% in simulation. Thus, our planning tools perform with limited success when there is no overspray pass. Nonetheless, our procedures yield desired uniformity results (when sufficient overspray is available) in a reasonable amount of time, and can offer guidelines to the paint specialists to substantially reduce the robot programming time.

6 Conclusion and Future Work

In this work, our approach examines the relationship between the variables and the output characteristics of the uniform coverage problem on simple yet realistic surfaces, and identifies the underlying key issues. We then decompose the coverage problem into three relatively independent sub-problems, thus significantly reducing the huge dimensionality of the search space for the coverage problem. This dimensionality reduction enables us to solve the coverage problem in a reasonable amount of time without over-trivializing assumptions about the deposition pattern or the surface geometry. Although we consider the deposition models only for ESRB atomizers, our coverage procedure is independent of the deposition model and can be used for any other type of deposition patterns.

The simulation results show that in presence of a sufficient overspray, the sub-problems of speed optimization along the passes and the index width optimization stay relatively independent. In the absence of sufficient overspray, deposition constraints in speed optimization problem may affect the index optimization results adversely, thus warranting a reformulation of the speed optimization and index width optimization problems. Our future work will address this reformulation of speed and index optimization subproblems, and will focus on developing path planning algorithms for broader class of surfaces, including surfaces with holes.

7 Acknowledgements

This work was supported by the National Science Foundation and the Ford Motor Company through grant IIS-9987972. The authors gratefully acknowledge the input

Target Surface	Norm. std. dev. (%)
Planar Sheet	2.64%
100cm Radius Cylinder passes parallel to axis	2.21%
100cm Radius Cylinder passes around cylinder	4.19%
Approximation to Ford Excursion Door Passes along zero curvature direction	6.81%
Approximation to Ford Excursion Door Passes along max. curvature direction	4.48%
Approx. to Crown Victoria fender	11.92%

Table 2. Index Width Optimization Results

and support provided by Dr. Jake Braslaw and Mr. Jeff Petty, our collaborators at the Ford Motor Company. We also sincerely thank Mr. Tie Wang from the Ford Motor Company and Mr. Greg Marsh, Mr. Noel Gauci, Mr. Michael Bichel, Mr. Ron Mata and Mr. Dale Smith from the ABB for the help provided in carrying out the experiments at the ABB facility in Auburn Hills, Michigan, USA.

References

1. Suk-Hwan Suh, In-Kee Woo, and Sung-Kee Noh. Development of An Automated Trajectory Planning System (ATPS) for Spray Painting Robots. In *IEEE Int'l. Conf. on Robotics and Automation*, pages 1948–55, Sacramento, California, USA, April 1991.
2. M. A. Sahir and Tuna Balkan. Process Modeling, Simulation, and Paint Thickness Measurement for Robotic Spray Painting. *Journal of Robotic Systems*, Vol. 17(9):479–94, 2000.
3. Naoki Asakawa and Yoshimi Takeuchi. Teachless Spray-Painting of Sculptured Surface by an Industrial Robot. In *IEEE Int'l. Conf. on Robotics and Automation*, volume 3, pages 1875–79, Albuquerque, New Mexico, USA, April 1997.
4. Weihua Sheng, Ning Xi, Mumin Song, Yifan Chen, and Perry MacNeille. Automated CAD-Guided Robot Path Planning for Spray Painting of Compound Surfaces. In *IEEE/RSJ Int'l. Conf. on Intelligent Robots and Systems*, volume 3, pages 1918–23, 2000.
5. Wesley H. Huang. Optimal Line-sweep-based Decompositions for Coverage Algorithms. In *IEEE Int'l. Conf. on Robotics and Automation*, volume 1, pages 27–32, Seoul, Korea, May 2001.
6. Taejung Kim and S.E.Sarma. Optimal Sweeping Paths on a 2-Manifold: A New Class of Optimization Problems Defined by Path Structures. *IEEE Transactions on Robotics and Automation*, 19(4):613–636, August 2003.
7. Tait S. Smith, Rida T. Farouki, Mohammad al Kandari, and Helmut Pottman. Optimal Slicing of free-form surfaces. *Computer Aided Geometric Design*, 19:43–64, 2002.
8. Ramanujam Ramabhadran and John K. Antonio. Fast Solution Techniques for a Class of Optimal Trajectory Planning Problems with Applications to Automated Spray Coating. *IEEE Transactions on Robotics and Automation*, Vol. 13(4):519–30, August 1997.
9. Eckhard Freund, Dirk Rokossa, and Jürgen Roßmann. Process-Oriented Approach to an Efficient Off-line Programming of Industrial Robots. In *IECON '98: Proceedings of the 24th Annual Conference of the IEEE Industrial Electronics Society*, volume 1, pages 208–13, 1998.
10. Heping Chen, Weihua Sheng, Ning Xi, Mumin Song, and Yifan Chen. Automated Robot Trajectory Planning for Spray Painting of Free-Form Surfaces in Automotive Manufacturing. In *IEEE Int'l. Conf. on Robotics and Automation*, volume 1, pages 450–55, Washington, D.C., USA, May 2002.
11. David C. Conner, Prasad N. Atkar, Alfred A. Rizzi, and Howie Choset. Development of Deposition Models for Paint Application on Surfaces Embedded in \mathbb{R}^3 for Use in Automated Path Planning. In *IEEE/RSJ Int'l. Conf. on Intelligent Robots and Systems*, volume 2, pages 1844–49, Lausanne, Switzerland, 2002.
12. David C. Conner, Prasad N. Atkar, Alfred A. Rizzi, and Howie Choset. Deposition Modeling for Paint Application on Surfaces Embedded in \mathbb{R}^3 for use in Automated Trajectory Planning. Technical Report CMU-RI-TR-02-08, Carnegie Mellon, Robotics Institute, Pittsburgh, Pennsylvania, USA, June 2002.
13. Prasad N. Atkar, Howie Choset, and Alfred A. Rizzi. Towards Optimal Coverage of 2-Dimensional Surfaces Embedded \mathbb{R}^3 : Choice of Start Curve. In *IEEE/RSJ Int'l. Conf. on Intelligent Robots and Systems*, volume 3, pages 3581–87, Las Vegas, USA, 2003.
14. Thomas Rausch, Franz-Erich Wolter, and Oliver Sniehotta. Computation of Medial Curves on Surfaces. *The Mathematics of Surfaces VII*, Information Geometers, 1997.