Maximizing Smart Factory Systems by Incrementally Updating Point Clouds

Evan Shellshear, Rolf Berlin, and Johan S. Carlson ■ *Fraunhofer-Chalmers Center for Industrial Mathematics, Sweden*

Up-to-date and physically accurate models of manufacturing facilities can help decision makers assess existing assembly lines and evaluate the flexibility of their industrial assets. Incremental laser scanning of structural changes to a factory floor can help ensure models are up to date and avoid excessive disruptions to production.

The end of the 20th century and the start of the 21st century were marked by a fundamental change in consumers' perceptions of the products they buy. The one-size-fits-all strategy that allowed companies to exploit economies of scale and ease of manufacturing has become a disadvantage as customers demand more personalized products. The integration of customers and business partners in the business and value processes is now seen as essential.[1]

To achieve these goals (and others), many industry leaders have recognized that they need to reinvent the way their factories operate. This has led to the Industrie 4.0[2] and Industrial Internet initiatives,[3] which are based on cyber-physical systems and/or the Internet of Things. The use of cyber-physical systems to provide customers with personalized products is one of the major goals of Industrie 4.0.

Although there is a strong desire to offer consumers the most customization possible, this should be done in a way that makes sense.[1] Jianxin Jiao and Mitchell Tseng stated that a good guiding principle to meet consumers' desires is "to optimize external variety while keeping control over internal complexity" so that production strategy and profitability are never at stake.[4] Our general goal is to follow this idea by helping factory owners identify the processes that would benefit most from Industrie 4.0 investments and determine which processes already contain the desired flexibility for model changes. This then allows companies to devote resources to those cyber-physical systems that will have the largest impact, hence raising the chance of greater industrial adoption of Industrie 4.0 paradigms due to more successful implementations.

In this article, we examine an example of this type of decision process. We look at a rust treatment process that Volvo wanted to examine in order to assess its reusability for different car-chassis models. To solve this problem, we proposed using a data-driven approach based on the idea of a digital factory[5]—that is, a virtual representation of the production plants and processes that can facilitate virtual analysis. Our data-driven approach has the potential to offer significant savings[5] because it avoids the construction of physical pilot plants anddoes not disturb ongoing production. The techniques we present here allowed Volvo to avoid the previous slow and expensive physical verification. Furthermore, through incremental point cloud updating, these techniques also allowed Volvo to easily keep an up-to-date and accurate virtual model for future analyses.

We see our tools as being general and able to answer important geometrical questions for numerous production processes. For example, how can existing assembly lines be virtually assessed for

 Published by the IEEE Computer Society

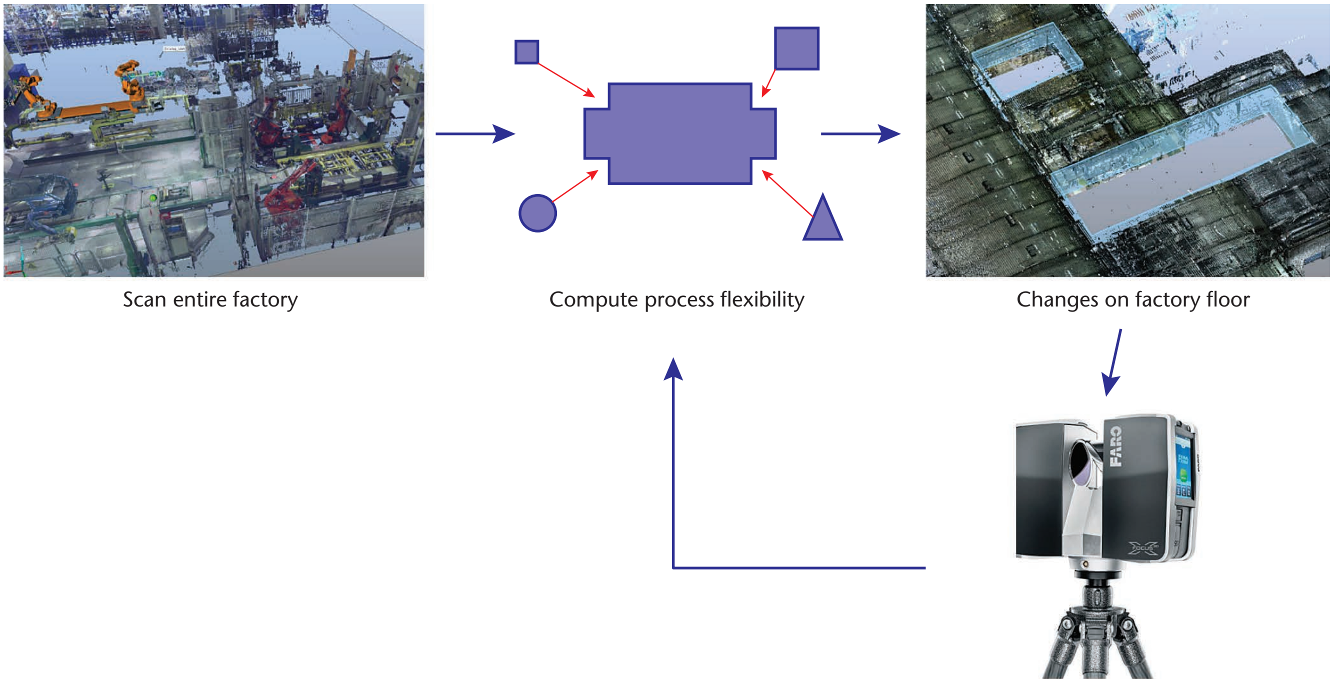

Figure 1. Process for determining physically accurate customization levels. After the initial laser scan of the entire manufacturing facility, it is possible to maintain up-to-date information with only single scans of small areas that usually take only 15 minutes per scan.

their maximum degree of customization? The general goal is to discover the maximum opportunity for customization without rebuilding or retooling existing facilities and, if rebuilding is required, to indicate the minimal necessary changes.

Our visualization tools realize the aforementioned goals by exploiting up-to-date information about the physical layout of the production environment. This is achieved via laser scan information that is collected and updated incrementally so that the designer's factory floor information is always accurate and up to date. This data maintenance is important because it has been recognized as a major stumbling block to the adoption of such models.[6] In addition, incrementally updating the laser scan data means that data collection is fast and efficient when compared with rescanning the entire area of interest. We also use this data in a computer model that allows for the flexibility of combining different sources of data (such as CAD and point cloud) to facilitate industrial-process simulations. This data is then used to compute maximum collision-free volumes that can move along the given path.

Our Process

Our process involves several interconnected steps that continually provide the designer with up-to-date information. Figure 1 illustrates the steps involved. The first step, the most time-consuming, involves the initial scanning of the manufacturing facilities.

After the initial facility-wide laser scan, maintaining the information so that it is up to date can be carried out efficiently and quickly. Updating the data involves only single scans of small areas, which usually take only about 15 minutes per scan. One way of doing this, which is in tune with the Industrie 4.0 revolution of continuous data collection for optimal adaptability, would be to have small permanent scanning installations scanning the factory area on a regular basis. In this study, the scanning was carried out manually. However, a permanent scanning installation would collect data that could then be immediately used to recompute the level of customization and confirm that the changes have not created any new limitations.

Factory Laser Scanning

During the last 10 years, 3D laser scanners have become a ubiquitous measurement technology. The laser scanners emit laser beams and compute the distances to objects by measuring the distance traveled by the laser beam (see Figure 2 for an example). For many factories, CAD data of the facilities is either missing or does not correspond to the final building.[7] For industrial settings, laser scanning technology can thus be used to provide accurate and cost-effective data for simulations.

Figure 2. A FARO 3D X 130 scanner used to scan industrial settings. The scanner emits laser beams and computes the distances to objects by measuring the distance traveled by the laser beam.

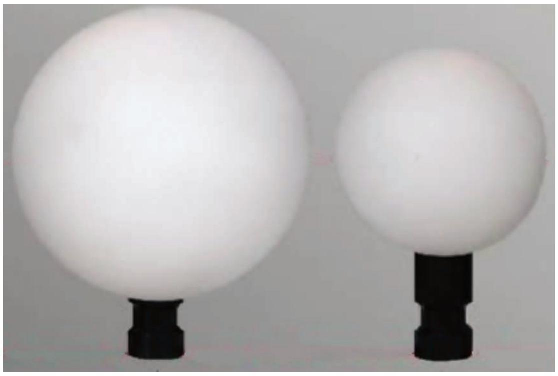

Figure 3. Spherical targets for registering scans. After all scans are complete, spherical markers are located in the point cloud data and used to register the scans. Thus, a single point cloud can be constructed from the data acquired from numerous spherical targets.

Although scanning technology may seem like a natural fit, most scanners have been developed for construction and surveying purposes, where one scans a scene hundreds of meters in size. Unfortunately, in production systems, the environment to be scanned is highly cluttered. Hence, for a similar area, the number of scans required can be significantly greater than a similar surveying scan. In the case study section, we present an example situation that took more than 50 scans (requiring two days and the collection of a total of 2 billion points) to acquire all the relevant detail, whereas a similar surveying scan over a comparable area would have been completed in one or two scans. Given the time required to acquire a complete scan, it is essential to avoid carrying out the entire process again to keep the data up to date. This motivation led to the development of the incremental updates technology.

For our intended use scenario, it is essential that decision makers are able to identify the existing flexibility in their production processes by having up-to-date and accurate information about their production processes. As we stated in the introduction, this information also needs to be acquired in a cost-efficient way that does not disturb current production strategies and processes. Although the initial scan of a factory can be expensive, incrementally updating the scanned point clouds can ensure the data is up to date in an inexpensive, fast way without disturbing the production process. Doing so lets production managers test new ideas (such as new cyber-physical systems) and quickly scan the new layout and approve all existing production processes with the new designs. This allows for fast iterations of customization and design changes without risking production halts resulting from unforeseen issues created in the design tests.

In this work, we used laser scanning because of its industrial maturity. For more information on other techniques that achieve a similar result (but that are more experimental or still at the research stage), see related research.[8]

Initial, Factory-Wide Scanning

Scanning a large facility such as a factory can require hundreds of individual data collections that then must be combined together to form a large single-point cloud. To combine this data, each scan needs to be located relative to the others. This can be achieved by numerous methods, but one typical method uses spherical targets[8] (see Figure 3). Once all necessary scans have been completed, these spherical markers are then located in the point cloud data and easily used to register the scans and create a single point cloud.[7]

To facilitate incremental updating of the point clouds, the spherical targets are attached to immovable objects in the factory (such as supporting columns). This is done so that the spherical target's mount can be left on the object, allowing easy incremental updates of the point cloud in the future. When the point cloud needs to be updated, the spherical target is attached to the mount and the area is rescanned. See Figure 4 for an example of a spherical mount.

Incremental Point Cloud Updates

Incremental scanning has become a key technique for updating point clouds during renovations or when retooling parts of factories. Scanning an entire factory can take days and disturb production, so incremental updates can help avoid this disruption for every small change.

To be able to incrementally update a point cloud, it is necessary to scan a factory in a well-defined *factory frame* (or coordinate system that is usually based on GPS coordinates) and reuse identical marker positions from previous scans.[8] If a section of a point cloud needs to be updated, we can select the relevant section of a point cloud by defining the area of interest's bounding box and then removing it. Because the new scan of the changed factory environment is scanned with the same factory frame, this removed section can be automatically added from the new scan to the old point cloud without any input from the user. Figure 5 shows an example of an incremental update of a section of a factory (the black volume).

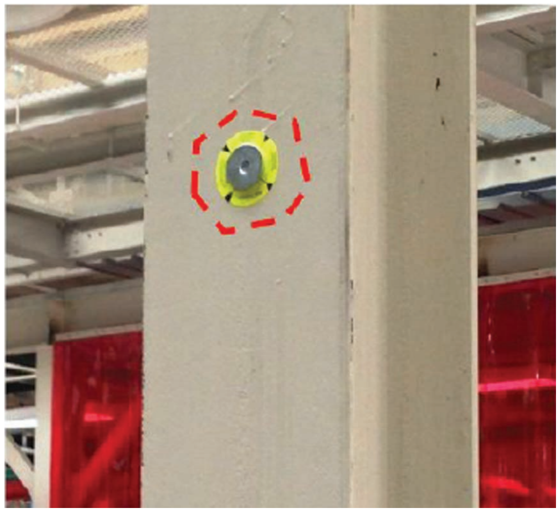

Figure 4. A mount (circled in red) for a spherical target. Mounts are left in place on immovable objects in the factory to facilitate future, incremental updates.

Computing the Customization Level

There are numerous factors that affect the level of customization of production processes; however, we focus on one critical aspect for which we have developed new visualization tools. The choice we make is based on the fact that many production processes are identical, regardless of an object's shape or size . For example, in the car industry, the processes for treating a chassis to protect it against rust and preparing a car for painting require that the entire chassis be sent through a number large basins filled with a special liquid. The limiting factor is simply the size and geometry of the chassis. The critical issue is whether a given part (such as the chassis) can actually pass from a one position to another without causing a collision.

This problem can be solved by simply computing whether new designs are collision free, but the most efficient way to answer this question is to provide the object's designer a maximum allowable volume. This can then be used to test new designs because any object that has a smaller volume and shape is guaranteed to function collision free with the given process. Using such techniques reduces the number of iterations by avoiding the testing of invalid (colliding) designs. In addition, such a volume can be a nice visual tool to assess design and customization constraints.

However, in scenarios where the object of interest cannot be changed and the surrounding infrastructure has to change, factory owners are often interested in the minimal factory changes possible to minimize downtime for certain processes. This problem can also be solved with the same framework used to solve the maximum-volume problem.

Computing the Maximum Volume and Colliding Infrastructure

There are two steps in providing the volume and collision information to the designer. The first step

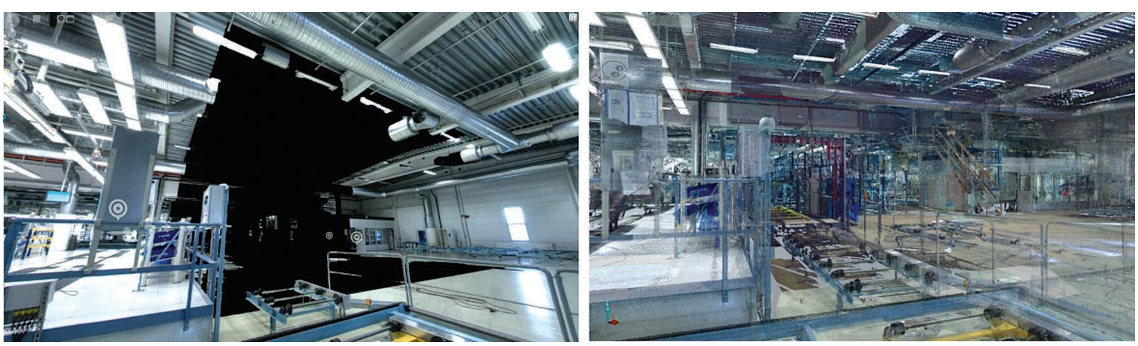

Figure 5. Incremental point cloud updates. The black volume indicates the section of the factory being updated.

Computing Collisions and Largest Volumes

In contrast to computing collisions, computing the largest volume that can move between two arbitrary points in space is a difficult problem. Here, we only require a solution to the simpler problem of computing the largest volume that can move along a fixed path, and this has been solved in two distinct ways. The difficulty is defining what is meant by the largest volume. Evan Shellshear and his colleagues solve this by voxelizing the volume of interest and defining the maximum volume simply to be a sum of all voxels.[1] Horea Ilies and Vadim Shapiro define the largest volume as the "dual of sweep" by finding the inverted motion of the object and computing the intersection of all positions for the inverted motion of the object.[2] The solution method proposed by both sets of authors involves the use of an Octree[3] to efficiently test for collisions between the surroundings and the moving object.

Both these methods have the advantage that the use of such a structure abstracts away the underlying geometries, so this structure can be used for point clouds, triangle meshes, quad meshes, NURBs, and so on. In addition, using such data structures makes it possible to specify a level of accuracy (Octree cell size) that can be used to take into account the object's elasticity or the oscillations along the path. Shellshear and his colleagues use the idea of representing both the moving object and static geometry as either an Octree or a set of voxels, and then they efficiently and correctly find collisions along the path using well-known results from path planning.[1] The methods from Ilies and Shapiro use the idea of a "point membership classification" combined with an Octree decomposition of space to efficiently find the largest volume contained in another object by finding points of collision.[2] The Shellshear methods were also shown to be fast on problems of practical interest and to scale well with the number of available processor cores.[1]

References

1. E. Shellshear, S. Tafuri, and J. Carlson, "A Multi-threaded Algorithm for Computing the Largest Non-colliding Moving Geometry," *Computer-Aided Design*, vol. 49, Apr. 2014, pp. 1–7.
2. H.T. Ilies and V. Shapiro, "The Dual of Sweep," *Computer-Aided Design*, vol. 31, no. 3, 1999, pp. 185–201.
3. D. Meagher, "Geometric Modeling Using Octree Encoding," *Computer Graphics and Image Processing*, vol. 19, no. 2, 1982, pp. 129–147.

is to compute the maximum volume that can pass along a given path without causing a collision. If there is no collision-free path, the designer will need to know the set of colliding points. There are several ways to do this;[9] however, our requirements are to be able to compute this information for massive point clouds containing hundreds of millions of points. We also want to be able to do this for incrementally updated point clouds. Although earlier methods were only designed for static point clouds,[7,9] they can be efficiently adapted for incremental updates as well.

For a point cloud with n points and a section to be updated with m points, we can update the data structures in both articles in $O(m \log n)$. Allowing process engineers to incrementally update the point cloud, instead of having to recompute all data structures from scratch, can save a significant amount of time and has numerous advantages. (In our case study, the update required only 15 minutes instead of two days.) Besides permitting process engineers to easily add the latest information about the factory floor to a computation, the incremental updating technique also can be used to quickly test virtual changes in a scanned version of a factory. This check can verify that no new conflicts have been introduced and also allow designers to see the effect of the changes on the largest volume object.

The second step in our process is to visualize the maximum volume and colliding infrastructure. Given the methods available for computing the largest volume (see the sidebar), the most natural representation of the result is to represent the final maximum-volume object as a set of (potentially different sized) boxes. Given the set of all noncolliding boxes, this information can be used to present the user with merely the outer boundary of the union of all noncolliding boxes, providing the user with a clean and easy-to-use representation of the maximum volume. However, a disadvantage with this representation occurs when the box size is small, which then creates a detailed representation of the maximum volume. Hence, methods to simplify the final representation are necessary so that boundary boxes defining a flat plane for a side are merged to minimize the number of triangles that are used to represent the box faces.

Figure 6 shows an example use of this tool for the path in Figure 7. The figure demonstrates a new car chassis model fitting inside the maximum volume that can move along a predefined path.

Computing the maximum volume (as described in the sidebar) also allows the designer to take into account unexpected movements (such as oscillations, elasticity, and vibrations) of the object . This is achieved by allowing the designer to specify a tolerance (a fixed distance that expresses the un-

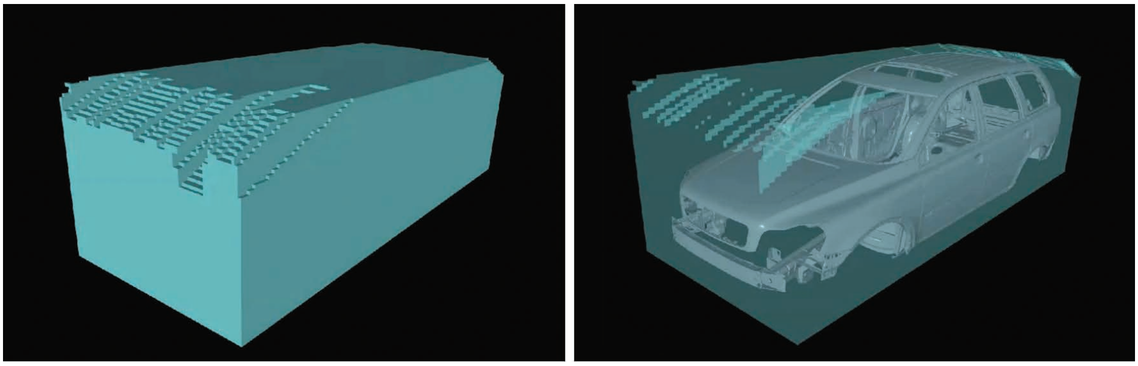

Figure 6. Visualizing the maximum volume and colliding infrastructure. The maximum collision-free volume for the path in Figure 7 with and without car chassis visible. (Courtesy of Volvo Cars.)

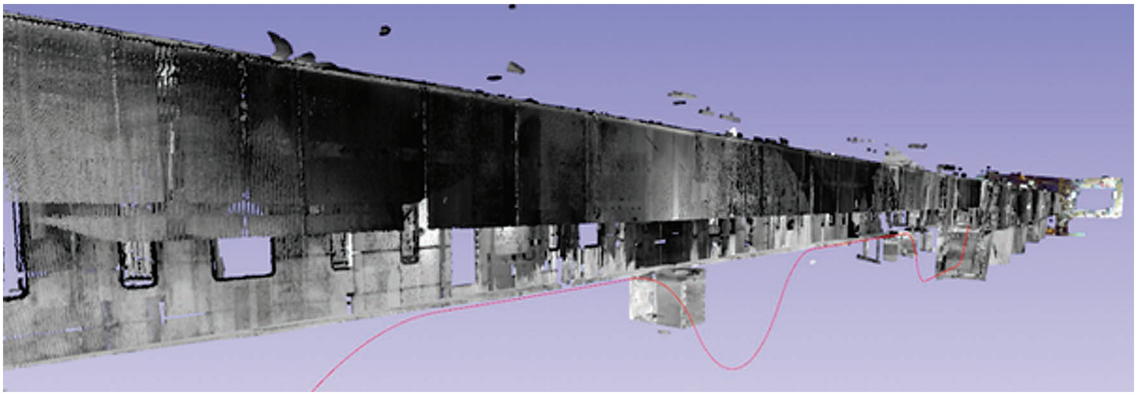

Figure 7. The scanned factory section. We computed the maximum volume (Figure 6) of the chassis moving along the path marked in red. A section of the point cloud has been removed so that the red path is visible. (Courtesy of Volvo Cars.)

certainty in the object's movement), which can then be translated into parameters defining the Octree (increasing or decreasing the minimum cell size and so forth).

When computing the set of colliding points for new objects moving along the path, the same methods can be reused. By voxelizing the geometry to a desired accuracy, we can then compute which points collide with the voxelized geometry along the path, and these colliding points can then be displayed. Figure 8 gives an example. Combining these two processes allows the designer to decide to either change the factory floor or change the design of the object. The tools presented here give the designer immediate feedback as to the feasibility of the intended changes.

Figure 8. Computing the set of colliding points for new objects. Example set of colliding points shown in red. (Courtesy of Volvo Cars.)

Case Study

A major impetus for the development of these tools has been to allow a larger variety of customized car chassis to take part in certain processes without redesigning the concerned parts of the factory. In particular, the rust protection process at Volvo that comprises a set of baths filled with liquids formed the basis for our investigations.

In the past, when designing a new model to pass through Volvo's set of rust protection baths, Volvo engineers couldn't be sure that the new model's chassis would actually fit along the path without causing collisions. To test whether a new design would fit, a number of expensive, inexact methods were developed.[6] For example, Volvo used to cut out a piece of cardboard representing the largest cross-section of the new chassis and then require someone to carry the piece of cardboard along the entire path (which can be hundreds of meters long). This required emptying the baths of liquid and was obviously an expensive, slow, and inaccurate process. On other occasions, even when there wasn't a new model to be tested, parts of the production process would require repairs that could affect the shape and form of the section that was being repaired. This would require the same cardboard cutout tests as in the previous case.

The tools presented in this article can improve such processes and give the car designer a powerful way to test car designs without requiring expensive and slow physical geometry verification. To leverage the tools developed here, the first step was to scan the entire rust protection area.

There are numerous laser scanners on the market, but the one used for the case study was a FARO 3D X 130 scanner (see Figure 2). The scanner has an accuracy of about ±2 mm at 100 meters and can scan both color and distance information. It takes approximately 5 to 7 minutes to complete a scan and approximately 15 minutes to set it up before scanning. Once the scanner is set up, the user simply presses a button and leaves the scanner to automatically record its surroundings.

The area to scan was 250 meters long and required 50 scans, resulting in a point cloud with approximately 2 billion points. Each scan took approximately 20 minutes, resulting in about two days of work. This point cloud was then filtered, and redundant points were removed via registration, resulting in 42 million points. Figure 7 shows a scan of the relevant part of the factory, with the path of the object in red. Given the point cloud, the tools from the previous section were used to assist the designer in knowing whether new designs could move along the rust protection section.

As we mentioned earlier, occasionally the rust protection part of the factory was updated, and by using incremental scans, we were able to reduce the time required for acquiring a physically correct model from two days to 20 minutes. This was done by combining the scanned part of the factory with the CAD data from the car chassis models to carry out simulations. We were also able to carry out the simulations quickly (within 10 minutes) due to the level of accuracy required. Because of the sideways movements of the vehicle along the path caused by side-to-side rocking of the chassis, accuracies of around 5 cm were acceptable (which translated into minimum Octree cell sizes of 5 cm). Such accuracies allowed for fast computations, taking less than 10 minutes, and hence fast design iterations. With the data from incremental updates, we easily updated the virtual factory model and were able to verify the compliance of all car models within an hour, saving days of manual verification.

The Industrie 4.0 movement presents companies with an unprecedented mindset to effectively answer consumers' demands for products more customized to their desires and needs. In spite of the possibilities envisioned by the Internet of Things and cyber-physical systems, companies will still be constrained by their production strategies and costs. The tools we have presented here can help factory owners better understand the level of customization already inherent in their current systems and make slight modifications as necessary. This has the important potential to clearly delineate the areas containing enough flexibility for future production models. Designers can then implement changes where they will provide the most benefit, before they embark on an expensive and potentially unnecessary reconstruction of an entire factory.

The ideas in this article could be taken further by augmenting the tools here with augmented reality. As discussed in earlier work,[10] the use of augmented reality would allow factory workers to overlay the information computed in this article with the real world and discuss the results with coworkers with the physical reality in front of them. Such a cyber-physical interface would provide an advantage due to the ability to both visualize the changes where they will occur and physically interact with the areas of interest.

The tools presented here significantly improved the process described in the case study section, but a number of challenges still remain. The first is that described in the "Computing Collisions and Largest Volumes" sidebar and occurs in situations where the path for a process is unknown. The choice of path affects the size of the largest volume, so it is important to have a way to choose a "good" path. Little research has been done to answer this challenge. In addition, it would be desirable to update changes to the factory floor in real time, although determining how to handle the data stream presents a major challenge.

Acknowledgments

We gratefully acknowledge Volvo Cars for the generous allowance for publication of the geometries presented here. This work was carried out at the Wingquist Laboratory VINN Excellence Centre and is part of the Sustainable Production Initiative and the Production Area of Advance at Chalmers University of Technology. It was supported by the Swedish Governmental Agency for Innovation Systems. We are also grateful for the kind help from ATS with data collection and analysis and are greatly indebted to the insightful *comments of the reviewers, who helped significantly improve the presentation and content of this article.*

References

1. B.H.M. Gerritsen, "Advances in Mass Customization and Adaptive Manufacturing," *Proc. Int'l Conf. Tools and Methods of Competitive Eng.* (TMCE 2008), vol. 2, 2008, pp. 869–880.
2. J. Jasperneite, "Was hinter Begriffen wie Industrie 4.0 steckt" [What Is behind Terms Such as Industrie 4.0], 19 Dec. 2012; www.computer-automation.de/steuerungsebene/steuern-regeln/artikel/93559/.
3. P.C. Evans and M. Annunziata, "Industrial Internet: Pushing the Boundaries of Minds and Machines," GE, 26 Nov. 2012; www.ge.com/docs/chapters/Industrial_Internet.pdf.
4. J. Jiao and M.M. Tseng, "Customizability Analysis in Design for Mass Customization," *Computer-Aided Design*, vol. 36, no. 8, 2004, pp. 745–757.
5. M. Gregor et al., "Digital Factory," *J. Automation Mobile Robotics and Intelligent Systems*, vol. 3, no. 3, 2009, pp. 123–132.
6. J. Svedberger and J. Andersson, "Laser Scanning in Manufacturing Industries," master's thesis, KTH Industrial Eng. and Management, 2013.
7. H.T. Ilies and V. Shapiro, "The Dual of Sweep," *Computer-Aided Design*, vol. 31, no. 3, 1999, pp. 185–201.
8. E. Lindskog et al., "Lean Based Problem Solving Using 3D Laser Scanned Visualizations of Production Systems," *Int'l J. Eng. Science and Innovative Technology*, vol. 3, no. 3, 2014, pp. 556–565.
9. E. Shellshear, S. Tafuri, and J. Carlson, "A Multi-threaded Algorithm for Computing the Largest Non-colliding Moving Geometry," *Computer-Aided Design*, vol. 49, Apr. 2014, pp. 1–7.
10. R. Schönfelder and D. Schmalstieg, "Augmented Reality for Industrial Building Acceptance," *Proc. IEEE Virtual Reality Conf.* (VR 2008), 2008, pp. 83–90.

Evan Shellshear *is an applied researcher at the Fraunhofer-Chalmers Center for Industrial Mathematics, Sweden. His research interests include computational geometry and computer graphics. Shellshear has a PhD in game theory from Bielefeld University. Contact him at evan.shellshear@fcc.chalmers.se.*

Rolf Berlin *is a CTO at ATS AB. His research interests include high-quality 3D data capturing. Berlin has a Tek Lic in mobile robotics from Chalmers University of Technology. Contact him at Rolf@ats.se.*

Johan S. Carlson *is the director of the Fraunhofer-Chalmers Centre. His research interests include virtual product realization and in particular geometry simulation and assurance. Carlson has a PhD in mathematical statistics from Chalmers University of Technology. Contact him at johan.carlson@fcc.chalmers.se.*