BEST PRACTICES IN MOVING TOWARDS LARGE-SCALE **PRODUCTION PLANNING IN 3D PRINTING**

Planning and scheduling for additive manufacturing (AM) at a factory scale opens up an incredible array of possibilities, but also corresponding challenges. At the top of those challenges is the fact additive manufacturing - or colloquially 3D printing - tends to get hung up on a single, time-consuming step.

The positive side of this bottleneck is that there is an opportunity to fit multiple components onto the same build plate, so they can be produced in parallel while incurring only marginal additional time costs. This paper is about finding ways to optimize the process and incorporate the best aspects of 3D printing and traditional manufacturing together to create a better future for everyone.

LOOKING BACK, LOOKING AHEAD

AM has been around in various forms for decades, at least since the invention of stereolithography, but the concepts behind it go back a century or more.

Over that time, there have been some wrong turns and unfulfilled promises, but in 2018, 3D printer tech is entering a new phase of mastery and optimization. Recent advances in AM have brought widescale adoption and refinement to 3D printing, allowing for the practical realization of designs that were impossible before.

Naturally, this next stage requires more intense scrutiny on methods for solving new optimization challenges. Fortunately, roadblocks to commercial 3D printing have attracted some of the brightest minds in fields as diverse as mechanical engineering, materials science, operations research, and production engineering. Any sustained optimization framework should incorporate the collaborative findings from all of these experts.

A 360 DEGREE REVIEW OF 3D PRINTING

The roots of AM technology date back to the 1800s with developments in the fields of photo sculpture and topography. However, the overall vision of printing three-dimensional objects was not really possible until advancements in materials science, computer-aided design (CAD), computer numerical control (CNC), and laser technology triggered the rapid development in AM techniques

There are seven major technology categories of AM today[1]:

- **Binder jetting** → a liquid bonding agent is selectively deposited to join powder materials
- **Directed energy deposition** → focused thermal energy is used to fuse materials by melting as they are deposited
- **Material extrusion** → material is selectively dispensed through a nozzle
- **Material jetting** → droplets of build material are selectively deposited on a surface
- **Sheet lamination** → sheets of material are bonded to form an object
- **Vat photopolymerization** → liquid photopolymer in a vat is selectively cured by light-activated polymerization
- **Powder bed fusion** → thermal energy selectively fuses regions of a powder bed

Powder bed fusion includes a special form of AM known as selective laser melting (SLM), which is typically used for producing metallic parts. Thanks to the immense industrial value of fusing metal powder into custom shapes across many different manufacturing verticals, SLM is the main focus of this article.

Schematic of selective laser melting (SLM), which is classified as a powder bed fusion process[2]

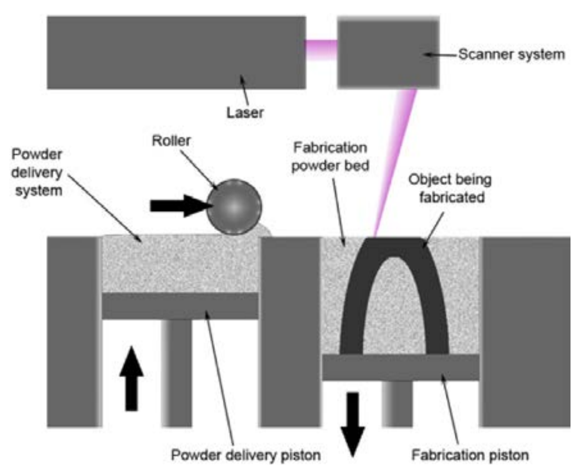

[1] ASTM International. 2018. *Standard Terminology for Additive Manufacturing Technologies*. ISO/ASTM 52900.

[2] Wikimedia Commons. 2018. https://upload.wikimedia.org/wikipedia/commons/3/33/Selective laser melting system schematic.jpg. User:Materialgeeza / CC-BY-SA-3.0.

WHAT 3D PRINTING MEANS FOR ENGINEERS AND DESIGNERS

3D printing has solved problems beyond the simply practical. It has opened up a new world in terms of functionality and aesthetics as well. Parts have to work, but they also have to meet an engineer's sense of parsimony – elegance in design. As engineered components grow more advanced and intricate, they push the boundaries of what is manufacturable using the traditional processes.

Some of these boundaries are technical, such as when certain geometries or materials simply don't work on traditional equipment. More often, they are possible but don't make sense economically. The mix of high touch production with low volume of production orders ends up being a resource sink. However, these types of challenges are ideal for 3D printers. Now it's cost-efficient to create these parts in a new way that was impossible before due to the high tooling/labor costs of each new setup.

AM has presented new capabilities to engineers and designers, greatly expanding the domain of the manufacturable design space in a variety of ways, allowing for design behaviors and concepts that were impractical dreams before now.

3D printers have made real the former oxymoron: mass customization. Today even intricate geometries and a batch size of one are both economically feasible and within reach, with no restrictions on visibility, draft angles, or other design requirements. The most creative designers are taking advantage of exotic material behaviors such as negative stiffness and negative thermal expansion[3].

For the manufacturing and production engineers, AM technologies allow components to be printed "on demand", and for the supply chain to consist of digital files rather than physical components.

While the unit production cost for a single additive component may be higher when compared to full-scale traditional production, AM is drastically less expensive at low and medium volume production. The industry is beginning to embrace these new manufacturing methods in order to streamline aspects of production, such as the well-known example of GE printing fuel injector nozzles for jet engines in a single print[4]. The nozzle previously required producing of 20 complicated components, as well as the additional labor to weld and braze them into an assembly afterward. In addition to being more streamlined to produce, the additive version of the nozzle was also 25% lighter, which directly translates to additional fuel savings for the aircraft[5].

AM can result in drastic reductions in the overall production time of a component. This efficiency stems from reducing the number of procedures involved in the production process, even if a typical print takes longer than a typical process step in a traditional manufacturing flow.

---

[3] Duoss, E. B.; Weisgraber, T.; Hearon, K.; Zhu, C.; Small, W.; R. Metz, T.; Vericella, J.; D. Barth, H.; Kuntz, J.; Maxwell, R.; M. Spadaccini, C.; and Wilson, T. 2014. Three-dimensional printing of elastomeric, cellular architectures with negative stiffness. Advanced Functional Materials 24(31):4905–4913.

[4] Coykendall, J.; Cotteleer, M.; Holdowsky, J.; and Mahto, M. 2015. 3D opportunity in aerospace and defense: additive manufacturing takes flight. *A Deloitte series on additive manufacturing.*

[5] Conner, B. P.; Manogharan, G. P.; Martof, A. N.; Rodomsky, L. M.; Rodomsky, C. M.; Jordan, D. C.; and Limperos, J. W. 2014. Making sense of 3-D printing: Creating a map of additive manufacturing products and services. *Additive Manufacturing* 1-4:64–76.

CONTROL MECHANISMS

Other challenges related to engineering involve understanding, modeling, and controlling complex physical behaviors and phenomena involved in the process, which are ultimately influenced by the planning and scheduling decisions made upstream of the process.

Some examples of these decisions are how to optimally orient the geometries in space for printing, and how to optimally nest multiple parts within a single print without inducing any build failures. “Optimal” may be defined in terms of overall print cost, overall print quality, or some other metric of optimality. It is important to note that some process decisions may induce physical conditions within the print which yield total process failures, with the onset of cracking, layer delamination, or unacceptable degrees of component warping. For this paper, we will assume that planning and scheduling decisions are made in fully observable and deterministic world where all possibilities succeed.

Additive manufacturing has been an emerging technology for several decades, and each advancement has provided new alternatives to traditional manufacturing methods for candidate geometries and applications. Mass production via additive manufacturing is now viable, yet the operations research techniques designed for traditional manufacturing are not directly transferable to AM. It’s time for new optimization models that take the latest advances into consideration.

OPTIMIZING 3D PRINTING FOR THE MARKET

Consider the following business use case for SLM. A company must produce a set of parts with unique configurations under a tight deadline. The company has to think about how to improve 3D printing and how it fits into their entire production process.

While the 3D printers have to be set up to minimize production time, satisfying deadlines also involves bringing together bin packing, nesting (two-dimensional bin packing), job shop scheduling, and the diverse requirements of all the company’s stakeholders.

Production planning in AM starts when a customer sends design specifications (CAD file) of a part that needs to be manufactured, along with a production deadline and delivery location. The design specifications of the part include geometric dimensioning and tolerancing (GD&T) information, which restricts the decision space in the production planning process. For example, the GD&T information may specify tolerances in regions of the design which are not achievable by some types of printing processes. As 3D printers handle multiple parts simultaneously in a single print operation, it is important to optimize the constitutive members of such sets and consider how they fit together.

Each part has its printing orientation, and they must collectively fit into the build volume of the machine – typically a regular hexahedron defined by height, width, and length. A single print operation is called a build, and we assume that the downward projections of different parts within a build cannot overlap. The duration of each build is functionally dependent on the speed that the laser beam in the machine is programmed to travel, the maximum height of the set of parts in the build, and the total path traveled by the laser beam during the build.

The problem consists of a set of parts, N, including their possible orientations, configurations, and deadlines, as well as a set of machines, M. The solution to the problem is the set of builds, B, scheduled to machines, M, which maximizes the number of parts printed prior to their deadlines, while also minimizing the overall printing cost.

In this case, you can disregard requirements normally considered during production planning. In particular, the shipping deadlines to given delivery locations are simplified by projecting the upper bound of shipping time into the scheduling deadline.

In the real world, the location of manufacturing assets around the globe would influence scheduling deadlines, but upper bounding of resource requirements for logistics allows for treating all machines as though they are in a single location.

Also, it makes sense to consider 2D bin-packing (nesting) when combining different parts into a single build. While 3D bin-packing is conceptually possible, i.e. by using horizontal bridges, all of that is outside the bounds of this article.

Finally, builds may require additional preparation and post-processing steps when the machine is cleaned, the material is recycled, the machine configuration is adjusted, and small finishing operations are required on the parts before shipment.

Those steps are performed by human operators that have their own shifts and schedules that vary across different time-zones (i.e. a high priority order may start to be printed in Hungary, because the operators are already up and can configure the machines, while a factory in Nevada is printing their night jobs). We are relaxing the human operators and upper bounding their work into the duration of printing a build, which also allows for relaxing the temporal reasoning to the sequence of builds on each machine.

The problem we are solving consists of multiple nested NP-hard problems (bin-packing, set cover, job-shop), and to describe the structure of our model we first encapsulate the configuration of parts. In the following sections we start by describing how the parts are nested together in builds. Then we provide a mathematical model consisting of variables, constraints and an objective function. Finally, we show and discuss experimental results.

NESTING

The nesting algorithm aims to efficiently arrange a given set of parts on a given build plate size by maximizing the total area covered with parts. Each part has an associated economic value and several possible orientations, which are considered under several (< 10) different rotations around the z-axis. The nesting algorithm maximizes the value of the build plate.

Nesting operates at three levels of fidelity with varying degrees of computational complexity - the highest fidelity nesting takes the longest to compute, while the lowest fidelity nesting is computed the fastest. Low-fidelity nesting uses 2D bounding boxes, medium-fidelity nesting uses silhouettes while packing the parts greedily, and high-fidelity nesting explores the search space using several search algorithms.

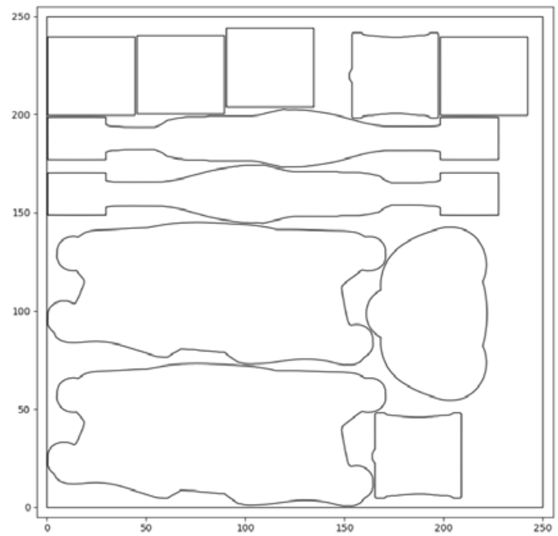

Example result from medium fidelity nesting.

MATHEMATICAL MODEL

We build up a representation of the problem as a meta constraint optimization problem[6] (V, C), where V is a set of variables and C is a set of constraints on top of V, where some constraints are hard and always need to be satisfied (a single machine cannot perform two activities at once), while others are soft (not every part can always be finished on time).

Then the solution of the problem (V, C) is an assignment that assigns exactly one value to each variable while all the hard constraints are satisfied. The quality of the solution is further evaluated based on the number of unsatisfied constraints, which is represented in the objective function. On top of the variables we are going to use multiple constraints and a strategy for ordering builds at the individual machines, described in the following sections.

[6] Dechter, R. 2003. Constraint Processing . San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

CONSTRAINTS

We use a number of standard constraints to maintain the consistency of the solution and cut symmetrical subsets from the search space. Additionally, we have developed three global constraints for additive manufacturing.

Compatibility Constraints enforce that parts within a single build are compatible with themselves, as well as with the machine which is processing the build. The compatibility between parts and machines captures an existence of a configuration that satisfies the configuration requirements from all the parts and that the machine that processes the parts supports given configuration.

Nesting Constraints are global constraints whose evaluation runs asynchronously in parallel with the main search algorithm. The main search algorithm uses the total surface area as an upper bound on how densely nested the parts can be in a build. Then those builds are sent to the nesting algorithm which provides a list of the parts that do fit into the build and the parts which are left over from the build. The constraint guarantees that all builds need to be empty or confirmed by nester. We treat the nesting constraint as a soft constraint, hence solutions with less violated nesting constraints are considered to be better quality.

Deadline Constraints guarantee that the parts are finished on time. We say that build's deadline is the earliest deadline among the deadline of its parts. Then, on a single machine the deadlines of all the builds must occur after the sum of durations of the builds whose deadlines are earlier. It may not be always possible to fulfill all the deadline constraints, and we may instead treat them as violated soft constraints which need to be minimized.

TARDY BUILDS

Once a build is assigned to a machine we can consider it to be independent of the builds assigned to the other machines. Such independence would disappear if we also considered that operators can service only a single machine at a single moment and whose availability is restricted. The advantage of the independence between machines is that instead of extending the state space with precedence constraints between builds on each machine we can choose a strategy that orders the builds.

Having multiple builds assigned to a single machine, we are mostly interested in whether a build is either on time, or if it is late and by how much it has missed its deadline.

We use an algorithm[7] optimal in minimizing the number of builds that have at least one order miss its deadline, however it does not guarantee optimality in minimizing the total number of orders that are late. In the trivial case, when each build consists of a single order, the algorithm is optimal for minimizing the number of late parts as well. However, minimizing the number of late parts is NP-hard in general.

We can show that when the deadlines for all parts are the same then the number of items in a build corresponds to a cost of an item in the knapsack problem. Consequently, we can translate a knapsack problem into minimizing the cost of parts that miss deadline.

[7] Pinedo, M. L. 2008. Scheduling: Theory, Algorithms, and Systems . Springer Publishing Company, Incorporated, 3rd edition.

OBJECTIVE FUNCTION

The primary optimization criterion is to minimize the number of unsatisfied soft (deadline and nesting) constraints, and then minimize the makespan (maximum execution time) of the schedule. Given that the nester is run asynchronously, its output becomes integrated into the main search algorithm via the lazy evaluation of the Nesting Constraints.

SOLVING APPROACH

Combining the previously defined structure and given a set of parts P and a set of machines M, we find the solution using following steps.

1. Create an initial state where each part is assigned to a different build and all the builds are assigned to some machines.
2. Perform a local search over the decision variables of the problem using the moves to change their values until the given time is spent.
3. Run nesting in parallel with the previous step, such that the builds confirmed by the nester are available in the nesting constraint and the nester is invoked whenever a new undiscovered and non-dominated build is considered by the nesting constraint.
4. Collect the assignments of parts to builds and schedule the builds to machines in time.

The simultaneous execution of steps 3 and 4 follows a best effort approach that allows two solvers for hard problems to exchange the information and reach higher quality solutions than if run in a sequence.

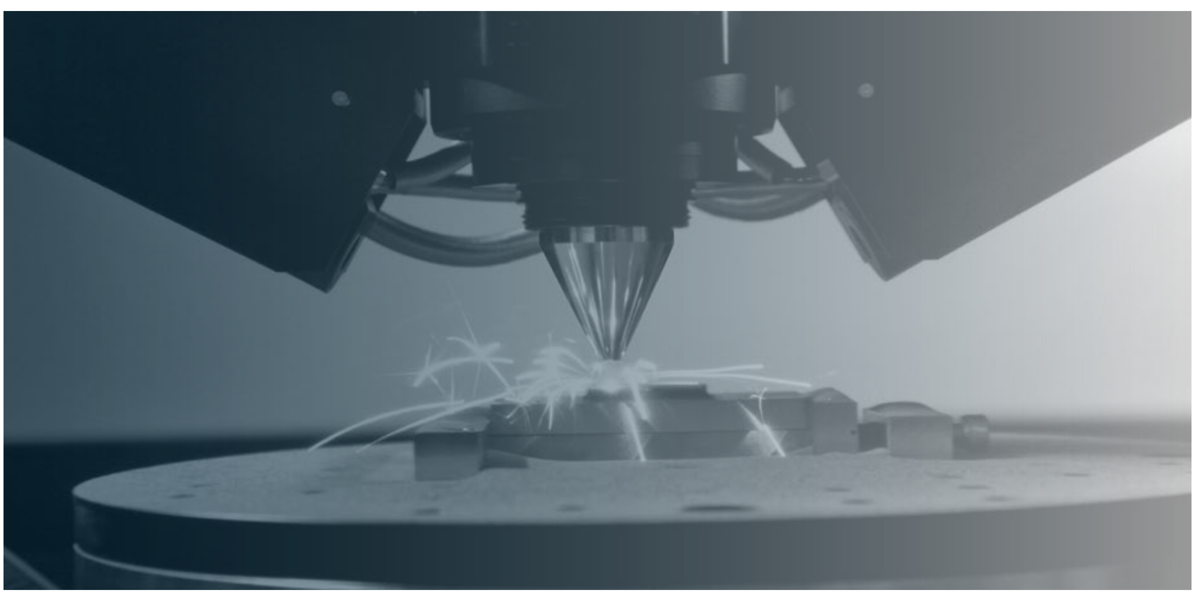

PRELIMINARY RESULTS

Using a generated dataset of 10 problem instances ranging from 150 to 1788 parts. We have run each search algorithm for 60 seconds and took out the best solution found. All of the algorithms found solutions which satisfy all of the deadline constraints, at which moment the optimization turns into makespan minimization.

The results indicate that even with an enormous search space it is possible to find reasonable solutions in a matter of minutes using simple local search algorithms and underlying constraint representations. Hill Climbing in particular seems to provide consistent performance, although it is sometimes surpassed by its variants that provide better makespans.

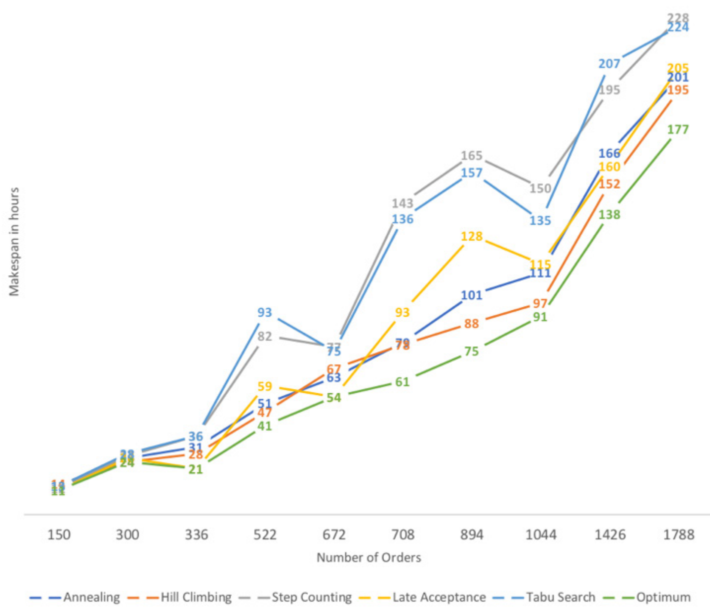

Shows the actual makespan in hours for how long it takes to execute the whole schedule. The graph includes the value for the precomputed optimal schedule and best schedule makespans found for other search algorithms run for 60 seconds.

Hill Climbing tends to get stuck in local optima traps and the other algorithms try to escape such traps through different relaxations of greediness of neighborhood exploration, leading to occasional successes such as Late Acceptance finding optimal solutions for 336 and 672 orders. The results are not conclusive with regard to which local search algorithm to choose, since the behavior can quickly change based on heuristics and neighborhood definitions, but it shows that local search is able to quickly provide reasonable solutions.

TAKEAWAYS

The focus of this article has been to capture the fundamental difficulty involved in the newly emerging intersection between AM, operations research, and AI. This model was presented at the 2018 International Conference on Automated Planning and Scheduling. We believe this new model we explored could ultimately be extended into an end-to-end optimization model that captures all the discrete decision-making challenges in AM.

AUTHORS

Filip Dvorak
Principal AI Engineer
Oqton, Inc.

Maxwell Micali
Research Engineer
Oqton, Inc.

Mathias Mathieu
3D Software Engineer
Oqton, Inc.

OQTON

San Francisco
832 Sansome Street, 4th Floor
San Francisco CA 94111
USA

Shanghai
29F, Jinmao Tower, 88 Century Ave
Shanghai 200121
China

Copenhagen
Lyngby Hovedgade 49A, 1 tv
2800 Kongens Lyngby
Denmark

Ghent
Sint-Jacobsnieuwstr 17, 2nd Floor
9000 Ghent
Belgium

www.oqton.com

info@oqton.com