

Co-Optimization of Design and Fabrication Plans for Carpentry

HAISEN ZHAO, University of Washington and Shandong University and IST Austria

MAX WILLSEY, AMY ZHU, CHANDRAKANA NANDI, and ZACHARY TATLOCK, University of Washington

JUSTIN SOLOMON, Massachusetts Institute of Technology

ADRIANA SCHULZ, University of Washington

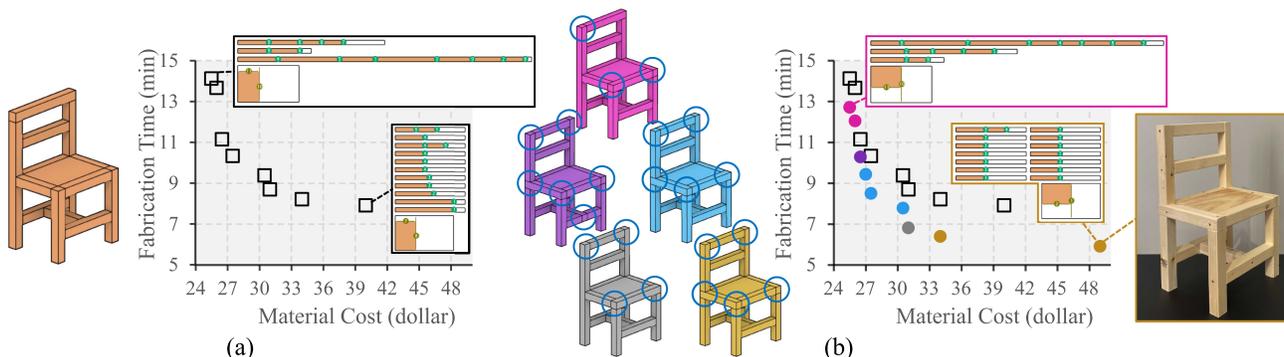


Fig. 1. Our system jointly explores the space of discrete design variants and fabrication plans to generate a Pareto front of (design, fabrication plan) pairs that minimize fabrication costs. In this figure, (a) is the input design for a chair and the Pareto front that only explores the space of fabrication plans for this design, (b) shows the Pareto front generated by joint exploration of both the design variants and fabrication plans for the chair, where each point is a (design, fabrication plan) pair. Design variations indicate different ways to compose the same 3D model from a collection of parts and are illustrated with the same color in the Pareto front. A physical chair is fabricated by following the result fabrication plan. The Pareto front of joint exploration dominates the Pareto front of (a), which shows that the fabrication cost can be significantly improved by exploring design variations.

Past work on optimizing fabrication plans given a carpentry design can provide Pareto-optimal plans trading off between material waste, fabrication time, precision, and other considerations. However, when developing fabrication plans, experts rarely restrict to a *single design*, instead considering *families of design variations*, sometimes adjusting designs to simplify fabrication. Jointly exploring the design and fabrication plan spaces for each design is intractable using current techniques. We present a new approach

to jointly optimize design and fabrication plans for carpentered objects. To make this bi-level optimization tractable, we adapt recent work from program synthesis based on equality graphs (e-graphs), which encode sets of equivalent programs. Our insight is that subproblems within our bi-level problem share significant substructures. By representing both designs and fabrication plans in a new *bag of parts* (BOP) e-graph, we amortize the cost of optimizing design components shared among multiple candidates. Even using BOP e-graphs, the optimization space grows quickly in practice. Hence, we also show how a feedback-guided search strategy dubbed *Iterative Contraction and Expansion on E-graphs* (ICEE) can keep the size of the e-graph manageable and direct the search towards promising candidates. We illustrate the advantages of our pipeline through examples from the carpentry domain.

CCS Concepts: • **Computing methodologies** → **Computer graphics**;

Additional Key Words and Phrases: Fabrication, programming languages

ACM Reference format:

Haisen Zhao, Max Willsey, Amy Zhu, Chandrakana Nandi, Zachary Tatlock, Justin Solomon, and Adriana Schulz. 2022. Co-Optimization of Design and Fabrication Plans for Carpentry. *ACM Trans. Graph.* 41, 3, Article 32 (March 2022), 13 pages.

<https://doi.org/10.1145/3508499>

1 INTRODUCTION

While optimizing designs for fabrication is a long-standing and well-studied engineering problem, the vast majority of the work in this area assumes that there is a unique map from a design to a fabrication plan. In reality, however, many applications allow for *multiple fabrication alternatives*. Consider, for example,

This material is based upon work supported by the National Science Foundation under Grant Nos. CCF-2017927, IIS-1954028, IIS-1838071, 1749570 and 1813166. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The work is also supported by the Google faculty award and the NSF China (62072284). Solomon acknowledges the generous support of Army Research Office grant W911NF2010168, of Air Force Office of Scientific Research award FA9550-19-1-031, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, from an MIT.nano Immersion Lab/NCSOFT Gaming Program seed grant, and from the Skoltech-MIT Next Generation Program.

Authors' addresses: H. Zhao, M. Willsey, A. Zhu, C. Nandi, Z. Tatlock, and A. Schulz, Paul G. Allen School of Computer Science & Engineering, University of Washington, 185 E Stevens Way, Box 352355, Seattle, WA, 98195-2355; emails: {haisen, mwillsey, amyzhu, cnandi, ztatlock, adriana}@cs.washington.edu; J. Solomon, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139; email: jsolomon@mit.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://www.acm.org/permissions).

© 2022 Association for Computing Machinery.

0730-0301/2022/03-ART32 \$15.00

<https://doi.org/10.1145/3508499>

the model is shown in Figure 1(a), where different fabrication plans trade-off material cost and fabrication time. In this context, fabrication-oriented design optimization becomes even more challenging, since it requires exploring the landscape of optimal fabrication plans for *many* design variations. Every variation of the original design (Figure 1(b)) determines a new landscape of fabrication plans with different cost trade-offs. Designers must therefore navigate the *joint* space of design and fabrication plans to find the optimal landscape of solutions.

In this work, we present a novel approach that simultaneously optimizes both the design and fabrication plans for carpentry. Prior work represents carpentry designs and fabrication plans as programs [Wu et al. 2019] to optimize the fabrication plan of a *single design* at a time. Our approach also uses a program-like representation, but we *jointly* optimize the design and the fabrication plan.

Our problem setting has two main challenges. First, the discrete space of fabrication plan alternatives can vary significantly for each discrete design variation. This setup can be understood as a *bi-level* problem, characterized by the existence of two optimization problems in which the constraint region of the upper-level problem (the joint space of designs and fabrication plans) is implicitly determined by the lower-level optimization problem (the space of feasible fabrication plans given a design). The second challenge is that there are multiple conflicting fabrication objectives. Plans that improve the total production time may waste more material or involve less precise cutting operations. Our goal is therefore to find *multiple* solutions to our fabrication problem that represent optimal points in the landscape of possible trade-offs, called the *Pareto front*. Importantly, the different fabrication plans on the Pareto front may come from different design variations. The complexity of the bi-level search space combined with the need for finding a landscape of Pareto-optimal solutions makes this optimization challenging.

We propose a method to make this problem computationally tractable in light of the challenges above. Our key observation is that there is redundancy on both levels of the search space that can be exploited. In particular, different design variations may share similar subsets of parts, which can use the same fabrication plans. We propose exploiting this sharing to encode a large number of design variations and their possible fabrication plans compactly. We use a data structure called an *equivalence graph (e-graph)* [Nelson 1980] to maximize sharing and thus amortize the cost of heavily optimizing part of a design since all other design variations share a part benefit from its optimization.

E-graphs have been growing in popularity in the programming languages community; they provide a compact representation for equivalent programs that can be leveraged for theorem proving and code optimization. There are two challenges in directly applying e-graphs to design optimization under fabrication variations, detailed below.

First, the different fabrication plans for a given design are all semantically equivalent programs. However, the fabrication plans associated with different design variations, in general, are *not semantically equivalent*, i.e., they may produce different sets of parts. This makes it difficult to directly apply traditional techniques, which exploit sharing by searching for minimal cost, but still semantically equivalent, versions of a program. One of our key tech-

nical contributions is therefore a new data structure for representing the search space, which we call the **Bag-of-Parts (BOP) E-graph**. This data structure takes advantage of common substructures across both design *and* fabrication plans to maximize redundancy and boost the expressive power of e-graphs.

Second, optimization techniques built around e-graphs have adopted a two-stage approach: *expansion* (incrementally growing the e-graph by including more equivalent programs¹) followed by *extraction* (the process of searching the e-graph for an optimal program). In particular, the expansion stage has not been feedback-directed, i.e., the cost of candidate programs has only been used in extraction, but that information has not been fed back in to guide further e-graph expansion. A key contribution of our work is a method for **Iterative Contraction and Expansion on E-graphs (ICEE)**. Because ICEE is feedback-directed, it enables us to effectively explore the large combinatorial space of designs and their corresponding fabrication plans. ICEE also uses feedback to prune the least valuable parts of the e-graph during the search, keeping its size manageable. Furthermore, these expansion and contraction decisions are driven by a multi-objective problem that enables finding a diverse set of points on the Pareto front.

We implemented our approach and compared it against prior work and against results generated by carpentry experts. Our results show that ICEE is up to 17 times faster than prior approaches while achieving similar results. In some cases, it is the only approach that successfully generates an optimal set of results due to its efficiency in exploring large design spaces. We showcase how our method can be applied to a variety of designs of different complexity and show how our method is advantageous in diverse contexts. For example, we achieve 25% reduced material in one model, 60% reduced time in another, and 20% saved total cost in a third when assuming a carpenter charges \$40/h when compared to a method that does not explore design variations.

2 RELATED WORK

Optimization for Design and Fabrication. Design for fabrication is an exciting area of research that aims at automatically achieving desired properties while optimizing fabrication plans. Examples of recent work include the computational design of glass façades [Gavriil et al. 2020], compliant mechanical systems [Tang et al. 2020], barcode embeddings [Maia et al. 2019], and interlocking assemblies [Cignoni et al. 2014; Hildebrand et al. 2013; Wang et al. 2019], among many others [Bickel et al. 2018; Schwartzburg and Pauly 2013]. Fabrication considerations are typically taken into account as constraints during design optimization, but these methods assume that there is an algorithm for generating *one* fabrication plan for a given design. To the best of our knowledge, no prior work explores the multi-objective space of fabrication alternatives during design optimization.

There is also significant literature on fabrication plan optimization for a *given* design under different constraints. Recent work includes optimization of composite molds for casting [Alderighi et al. 2019], tool paths for 3D printing [Etienne et al. 2019; Zhao et al. 2016], and decomposition for CNC milling [Mahdavi-Amiri et al. 2020; Yang et al. 2020]. While some of these methods minimize the

¹In the programming languages literature, this is known as *equality saturation*.

distance to a target design under fabrication constraints [Duenser et al. 2020; Zhang et al. 2019], none of them explores a space of design modification to minimize fabrication cost.

In contrast, our work *jointly* explores the design and fabrication space in the carpentry domain, searching for the Pareto-optimal design variations that minimize multiple fabrication costs.

Design and Fabrication for Carpentry. Carpentry is a well-studied domain in design and fabrication due to its wide application scope. Prior work has investigated interactive and optimization methods for carpentry design [Fu et al. 2015; Garg et al. 2016; Koo et al. 2014; Song et al. 2017; Umetani et al. 2012]. There is also a body of work on fabrication plan optimization [Koo et al. 2017; Lau et al. 2011; Leen et al. 2019; Yang et al. 2015]. Closest to our work is the system of Wu et al. [2019], which represents both carpentry designs and fabrication plans as programs and introduces a compiler that optimizes *fabrication* plans for a *single* design. While our work builds on the **domain specific languages (DSLs)** proposed in that prior work, ours is centered on the fundamental problem of design optimization under fabrication alternatives, which has not been previously addressed.

Bi-Level Multi-Objective Optimization. Our problem and others like it are *bi-level*, with a nested structure in which each design determines a different space of feasible fabrication plans. The greatest challenge in handling bi-level problems lies in the fact that the lower level problem determines the feasible space of the upper-level optimization problem. More background on bi-level optimization can be found in the book by Dempe [2018], as well as review articles by Lu et al. [2016] and Sinha et al. [2017].

Bi-level problems with multiple objectives can be even more challenging to solve [Dempe 2018]. Some specific cases are solved with classical approaches, such as numerical optimization [Eichfelder 2010] and the ϵ -constraint method [Shi and Xia 2001]. Heuristic-driven search techniques have been used to address bi-level multi-objective problems, such as genetic algorithms [Yin 2000] and particle swarm optimization [Halter and Mostaghim 2006]. These methods apply a heuristic search to both levels in a nested manner, searching over the upper level with NSGA-II operations, while evaluating each individual call in a low-level NSGA-II process [Deb and Sinha 2009]. Our ICEE framework also applies a genetic algorithm during the search. Different from past techniques, ICEE does not nest the two-level search but rather reuses structure between different upper-level feasible points. ICEE jointly explores both the design and fabrication spaces using the BOP E-graph representation.

E-graphs. An e-graph is an efficient data structure for compactly representing large sets of equivalent programs. E-graphs were originally developed for automated theorem proving [Nelson 1980], and were first adapted for program optimization by Joshi et al. [2002]. These ideas were further expanded to handle programs with loops and conditionals [Tate et al. 2009] and applied to a variety of domains for program optimization, synthesis, and equivalence checking [Nandi et al. 2020; Panckekha et al. 2015; Premtoon et al. 2020; Stepp et al. 2011; Wang et al. 2020; Willsey et al. 2021; Wu et al. 2019].

Recently, e-graphs have been used for optimizing designs [Nandi et al. 2020], and also for optimizing fabrication plans [Wu et al. 2019], but they have not been used to simultaneously optimize *both* designs and fabrication plans. Prior work also does not explore feedback-driven e-graph expansion and contraction for managing large optimization search spaces.

3 BACKGROUND

In this section, to increase the readability of the article and help readers get necessary background knowledge earlier, we introduce some mathematical preliminaries used in the rest of the article.

3.1 Multi-Objective Optimization

A multi-objective optimization problem is defined by a set of objectives $f_i : \mathbf{x} \mapsto \mathbb{R}$ that assign a real value to each point $\mathbf{x} \in \mathcal{X}$ in the feasible search space \mathcal{X} . We choose the convention that *small* values of $f_i(\mathbf{x})$ are desirable for objective f_i .

As these objectives are typically *conflicting*, our algorithm searches for a diverse set of points that represent optimal tradeoffs, called *Pareto optimal* [Deb 2014]:

Definition 3.1 (Pareto Optimality). A point $\mathbf{x} \in \mathcal{X}$ is *Pareto optimal* if there does not exist any $\mathbf{x}' \in \mathcal{X}$ so that $f_i(\mathbf{x}) \geq f_i(\mathbf{x}')$ for all i and $f_i(\mathbf{x}) > f_i(\mathbf{x}')$ for at least one i .

We use $F : \mathbf{x} \mapsto \mathbb{R}^N$ to denote the concatenation $(f_1(\mathbf{x}), \dots, f_N(\mathbf{x}))$. Pareto optimal points are the solution to the multi-objective optimization:

$$\min_{\mathbf{x}} F(\mathbf{x}) \quad \text{s.t. } \mathbf{x} \in \mathcal{X}. \quad (1)$$

The image of all Pareto-optimal points is called the *Pareto front*.

Non-Dominated Sorting. Genetic algorithms based on non-dominated sorting are a classic approach to multi-objective optimization [Deb and Jain 2013; Deb et al. 2002]. Sorting is the step of genetic algorithms that select parent populations for crossover and mutation. The key idea behind non-dominated sorting is that this selection should be done based on proximity to the Pareto front. These articles define the concept of Pareto layers, where layer 0 is the Pareto front, and layer l is the Pareto front that would result if all solutions from layers 0 to $l - 1$ are removed. When selecting parent populations or when pruning children populations, solutions in lower layers are added first, and when a layer can only be added partially, elements of this layer are chosen to increase diversity. Different variations of this method use different strategies for diversity; we use NSGA-III [Deb and Jain 2013] in our work.

Hypervolume. *Hypervolume* [Auger et al. 2009] is a metric commonly used to compare two sets of image points during Pareto front discovery. Intuitively, the hypervolume measures the region dominated by a given reference point and therefore a larger hypervolume implies a better approximation of the Pareto front. To calculate the hypervolume, we draw the smallest rectangular prism (axis-aligned, as per the L^1 norm) between some reference point and each point on the Pareto front. We then union the volume of each shape to calculate the hypervolume.

3.2 Bi-level Multi-Objective Optimization

Given a design space \mathcal{D} that defines possible variations of a carpentry model, our goal is to find a design $d \in \mathcal{D}$ and a corresponding fabrication plan $p \in \mathcal{P}^d$ that minimizes a vector of conflicting objectives, where \mathcal{P}^d is the space of fabrication plans corresponding to design d . This setup yields the following multi-objective optimization problem:

$$\min_{p,d} F(d,p) \text{ s.t. } d \in \mathcal{D}, p \in \mathcal{P}^d,$$

where \mathcal{P}^d defines the space of all possible plans for fabrication the design d . Generally, our problem can be expressed as a bi-level multi-objective optimization that searches across designs to find those with the best fabrication costs, and requires optimizing the fabrication for each design during this exploration [Lu et al. 2016]:

$$\min_d F(d,p) \text{ s.t. } d \in \mathcal{D}, p = \arg \min_p F(d,p),$$

where $\arg \min$ refers to Pareto-optimal solutions to the multi-objective optimization problem.

A naïve solution to this bi-level problem would be to search over the design space \mathcal{D} using a standard multi-objective optimization method, while solving the nested optimization problem to find the fabrication plans given a design at each iteration. Given the combinatorial nature of our domain, this would be prohibitively slow, which motivates our proposed solution.

3.3 Equivalence Graphs (E-graphs)

Typically, programs (often referred to as *terms*) are viewed as tree-like structures containing smaller sub-terms. For example, the term 3×2 has the operator \times at its “root” and two sub-terms, 3 and 2, each of which has no sub-terms. Terms can be expressed in multiple syntactically different ways. For example, in the language of arithmetic, the term 3×2 is semantically equivalent to $3 + 3$, but they are syntactically different. Naïvely computing and storing all semantically equivalent but syntactically different variants of a term requires exponential time and memory. For a large program, this makes searching the space of equivalent terms intractable.

E-graphs [Nelson 1980] are designed to address this challenge—an e-graph is a data structure that represents many equivalent terms efficiently by sharing sub-terms whenever possible. An e-graph not only stores a large set of terms, but it represents an *equivalence relation* over those terms, i.e., it partitions the set of terms into *equivalence classes*, or *e-classes*, each of which contains semantically equivalent but syntactically distinct terms. In Section 4.2, we show how to express carpentry designs in a way that captures the benefits of the e-graph.

Definition 3.2 (E-graph). An *e-graph* is a set of equivalence classes or *e-classes*. An *e-class* is a set of equivalent e-nodes. An *e-node* is an operator from the given language paired with some e-class children, i.e., $f(c_1, \dots, c_n)$ is an e-node where f is an operator and each c_i is an e-class that is a child of this e-node. An e-node may have no children, in which case we call it a *leaf*. An e-graph *represents* an equivalence relation over terms. Representation is defined recursively:

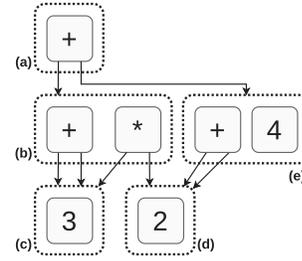


Fig. 2. An example e-graph. E-classes (dotted boxes labeled by letters) contain equivalent e-nodes (solid boxed) which refer to children e-classes (arrows). The e-class (c) contains one leaf e-node, 3, and it represents one term, 3. The e-class (b) contains two e-nodes, (c) + (c) and (c) * (d), and it represents two terms: $3 + 3$ and $3 * 2$. Although the e-class (a) only contains one e-node, it represents 4 terms: $(3 + 3) + (2 + 2)$, $(3 * 2) + (2 + 2)$, $(3 + 3) + 4$, and $(3 * 2) + 4$. If + is cheaper than *, then $(3 + 3) + 4$ is the cheapest term represented by e-class (a).

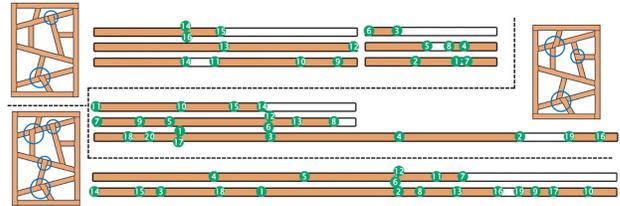


Fig. 3. Example of three different design variations of a model and corresponding fabrication plans. Design variations determine different ways to decompose a 3D model into a set of parts. Fabrication plans define how these parts are *arranged* in pieces of stock material and the cut order (illustrated by the numbers along with each cut).

- An e-graph represents a term if any of its e-classes do.
- An e-class represents a term if any of its e-nodes do. All terms represented by e-nodes in the same e-class are equivalent.
- An e-node $f(c_1, \dots, c_n)$ represents a term $f(t_1, \dots, t_n)$ if each e-class c_i represents term t_i . A leaf e-node g represents just that term g .

Figure 2 shows an example of an e-graph and representation. Note how the e-graph maximizes sharing even across syntactically distinct, semantically equivalent terms. When adding e-nodes or combining e-classes, the e-graph automatically maintains this maximal sharing property, using existing e-nodes whenever possible.

4 OPTIMIZATION ALGORITHM

Our algorithm takes as input a carpentry design with a discrete set \mathcal{D} of possible design variations. Design variations determine different ways to decompose a 3D model into a set of fabricable parts, as shown in Figures 3 and 4. These can be manually or automatically generated (see Section 1.1 of the supplemental material).

Our goal is to find Pareto-optimal solutions that minimize fabrication cost, where each solution is a pair of design variation and fabrication plan. Similar to prior work [Wu et al. 2019], we measure cost in terms of material usage (f_c), cutting precision (f_p), and fabrication time (f_t). Section 1.3 of the supplemental material describes how these metrics are computed for this work.

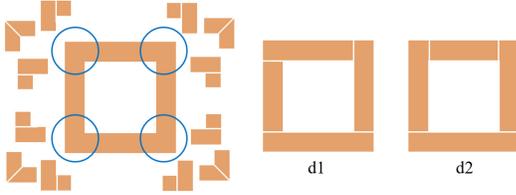


Fig. 4. Example of a space of design variations, \mathcal{D} . Each of the four connectors can have three different connecting variations, resulting in 81 design variations. Note that some of the different design variations may use the same parts (as d1, d2), and will be treated as redundant during our optimization. This model produces 13 unique bags of parts.

4.1 Motivation and Insights

Given an algorithm for finding the Pareto-optimal fabrication plans for a given design (e.g., the one proposed by Wu et al. [2019]), a brute force method would simply find the Pareto-optimal solutions for each of the possible design variations $d \in \mathcal{D}$ and take the dominant ones to form the Pareto front of the combined design/fabrication space. Since design variations can produce an exponentially large space of designs \mathcal{D} , this approach would be intractable for complex models. An alternative approach could use a discrete optimization algorithm to explore the design space (e.g., hill climbing). This approach would still need to compute the Pareto-optimal fabrication plans for each design explored in every iteration, which can be expensive for complex design variants (e.g., it takes 8–10 minutes to compute Pareto-optimal fabrication plans for a single design variation of the chair model in Figure 1 using the approach of Wu et al. [2019]).

We address these challenges with two key insights:

- (1) Design variants will share common sub-parts (both within a single variant and across different variants). As shown in Figure 3, even in a design where no two parts are the same, there is significant overlap *across* design variations. Exploiting this sharing can prevent recomputing the fabrication cost from scratch for every design variation. We propose using an e-graph to capture this sharing when (sub-)designs have the same BOP; we call this e-graph the *BOP E-graph*.
- (2) The space of design variants is too large to exhaustively explore, and even a single variant may have many Pareto-optimal fabrication plans. We propose using the BOP E-graph to guide the exploration in an incremental manner, with a new technique called *ICEE* that jointly explores the design and fabrication plan spaces.

4.2 Bag of Parts (BOP) E-graph

Our algorithm selects a Pareto-optimal set of fabrication plans, each of which will produce a design variation of the given model. A *fabrication plan* consists of four increasingly detailed things:

- (1) A **BOP**, a bag² (a.k.a. multiset) of atomic parts that compose the model.
- (2) An **assignment** that maps those parts to individual pieces of stock material (see the parts in stock pieces in Figure 3).

- (3) A **packing** for each piece of stock in the assignment that dictates *how* those parts are arranged in that stock (see the layout on the stock pieces in Figure 3).
- (4) A **cutting order** for each packing that specifies the order and the tool (chopsaw, tracksaw, etc.) used to cut the stock into the parts (see the numbers indicating the order in Figure 3).

We say that an **arrangement** is items 1–3: a BOP assigned to and packed within pieces of stock material, but *without cutting order decided*. We can create a language to describe arrangements; a term in the arrangement language is one of the following:

- An **atomic node** is a childless operator that represents a BOP packed into a single piece of stock. For example, $\{\square, \square, \Delta\}_{p,b}$ maps two squares and one triangle all to the same piece of stock of type b using a packing p .
- A **union node** takes two-child arrangements and composes them into a single arrangement. The following arrangement is a union node of two atomic nodes: $\{\square, \square\}_{p_1,b} \cup \{\Delta\}_{p_2,b}$. It packs two squares into the stock of type b using packing p_1 , and it packs a triangle into a *different piece* of stock of the same type b using packing p_2 .

To put arrangements into an e-graph, we must define the notion of equivalence that the e-graph uses to determine which e-nodes go in the same e-class. The more flexible this notion is (i.e., the larger the equivalence relation), the more sharing the e-graph can capture.

To maximize sharing, we say two arrangements are equivalent if they use the same BOP, even if those parts are assigned to different stock or packed differently. For example, $\{\square, \square\}_{p_1,b}$ is equivalent to $\{\square, \square\}_{p_2,c}$ even though they use different kinds of stock, and $\{\square, \square, \Delta\}_{p_3,b}$ is equivalent to $\{\square, \Delta\}_{p_4,b} \cup \{\square\}_{p_5,b}$ even though the former uses one piece of b stock and the latter uses two.

Given our arrangement language and the BOP notion of equivalence, we can now describe the central data structure of our algorithm, the *BOP E-graph*. Recall from Section 3.3 that e-nodes within an e-graph have e-class children rather than e-node children. So, viewing our arrangement language at the e-graph level, union e-nodes take two e-classes as children. All e-nodes in the same e-class are equivalent, i.e., they represent terms that use the same BOP but that arrange those parts differently into stock.

Figure 5 gives two example design variations and a BOP E-graph that captures a common sub-arrangement between the two. The e-classes E1 and E2 represent terms that correspond to the two box designs, and E4 captures ways to arrange the y and z parts which the variants share. The design variant including part w also captures sharing with itself: e-class E5 reuses the arrangement in e-class E9.

Note that arrangements and the BOP E-graph do not mention designs. We do not “store” designs in the e-graph, we just need to remember which e-classes represent bags of parts that correspond to designs that we are interested in. This can be done outside the e-graph with a mapping from designs to e-classes. Many designs (especially symmetric ones) may have the same BOP. We call an e-class that is associated with a design a *root e-class*, and we call a term represented by a root e-class a *root term*. The BOP E-graph itself does not handle root vs. non-root e-classes or terms differently, these are only used by the remainder of the algorithm to

²A *bag* or *multiset* is an unordered set with multiplicity, i.e. it may contain the same item multiple times. We will use the terms interchangeably.

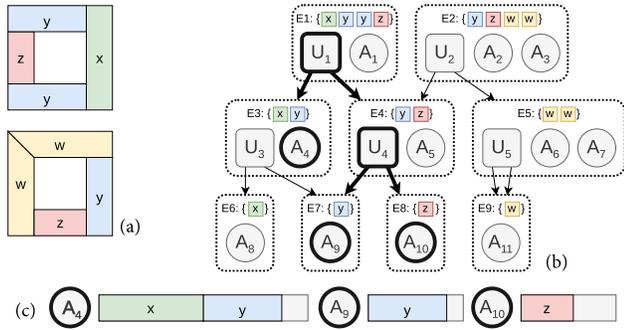


Fig. 5. Two variants (a) of a box design encoded in one BOP E-graph (b). The bold edges show a root term that requires three atomic packings (c). The BOP E-graph encodes multiple arrangements for both design variants. E-classes are drawn as dotted boxes and annotated with the BOP represented by that e-class. (Only the e-nodes are semantically *in* the e-graph; the name and BOP are just visual aides.) E-classes E1 and E2 are root e-classes since they represent the BOP required by the design variants. Union and atomic e-nodes are shown as squares with “U”s or circles with “A”s, respectively. Atomic e-nodes correspond to packings of parts within a piece of stock (c). An example root term in the BOP E-graph is bolded; using the syntax from Section 4.2, this is the term $\{x, y\}_{A_4, \text{long}} \cup \{y\}_{A_9, \text{short}} \cup \{z\}_{A_{10}, \text{short}}$.

remember which arrangements correspond to design variants. The BOP E-graph will maximize sharing across design variations *and* arrangements since it makes no distinction between the two.

4.3 Iterative Contraction and Extension on E-graphs (ICEE)

4.3.1 Overview. ICEE takes a feasible design space \mathcal{D} as input, and outputs a Pareto front where each solution s represents a (design, fabrication) pair. An overview of this algorithm is shown in Figure 6. The pseudocode is included in the supplemental material.

The initialization step selects a small subset of design variants from \mathcal{D} (Section 4.3.2) and then generates a small number of fabrication arrangements for each one (Section 4.3.3). All of these are added to the BOP E-graph, maintaining the property of maximal sharing, as described above. ICEE then applies the extraction algorithm (Section 4.3.4) to generate a Pareto front from the current BOP E-graph. This process will compute many different solutions s and their fabrication costs $F(s) = (f_m(s), f_p(s), f_t(s))$, all of which are stored in the solution set S .

The resulting Pareto front is used to compute ranking scores for each e-class in the BOP E-graph; the ranking score measures how often this BOP is used in Pareto-optimal solutions and how many fabrication variations have been explored for this BOP. Using these scores, ICEE contracts the BOP E-graph by pruning e-classes that have been sufficiently explored but still do not contribute to Pareto-optimal solutions (Section 4.3.5).

Having pruned the e-graph of the less relevant e-classes, ICEE then expands the BOP E-graph in two ways (Section 4.3.6). First, it suggests more design variations based on the extracted Pareto-Optimal designs. Second, it generates more fabrication arrangements for both the newly generated design variations and

some of the previously existing e-classes. The ranking scores are used to select e-classes for expansion.

ICEE then extracts the new Pareto front from the updated BOP E-graph and repeats the contraction and expansion steps until the following termination criteria are met: (1) There is no hypervolume improvement within t_d iterations, or (2) We exceed mt_d iterations. Additionally, we set a timeout T beyond which we no longer change the BOP E-graph, but continue to extract based on crossover and mutation until one of the termination criteria is met. In our experiments, we set $t_d = 10$, $mt_d = 200$, and $T = 4$ hours.

4.3.2 Initial Generation of Design Variants. We bootstrap our search with the observation that design variations with more identical parts tend to be cheaper to fabricate because less time is spent setting up fabrication processes. Therefore, instead of initializing the BOP E-graph with K_d designs randomly selected from \mathcal{D} , we randomly select up to 10^5 designs and select the top K_d designs from this set that have a maximal number of identical parts.

4.3.3 Fabrication Arrangements Generation. Again, instead of randomly generating K_f arrangement variations for a given design, we use heuristics; namely, that (1) We can minimize the number of cuts by stacking and aligning material to cut multiple parts with a single cut, and (2) We can minimize the material cost by packing as many parts as possible to a single stock. Since a similar method for generating arrangement variations has been previously proposed by Wu et al. [2019], we leave a detailed discussion of the algorithm for supplemental material (Section 1.2). We note that the key difference between our method and the prior heuristic-driven algorithm is that we incorporate storage and direct control schemes that enable the method to output K_f variations that are *different* from the ones generated during previous iterations of ICEE. This is essential to enable incremental expansion of the BOP E-graph without restoring variations that have already been pruned in the previous contraction steps.

4.3.4 Pareto Front Extraction. In e-graph parlance, *extraction* is the process of selecting the “best” represented term from an e-graph according to some (typically single-objective) cost functions. One way to view extraction is that it simply chooses which e-node should be the canonical representative of each e-class; once that is done, each e-class represents a single term. Since our cost function is multi-objective, we must instead extract a set of terms (arrangements) from the BOP E-graph that forms a Pareto front.

We use a genetic algorithm [Deb and Jain 2013] to extract terms from the BOP E-graph. The population size is set to N_{pop} . The genome is essentially a list of integers, one per e-class, that specifies which e-node is the representative. Since the BOP E-graph may have multiple root e-classes (corresponding to multiple design variations), we combine the genes for all the root e-classes, only picking a single e-node among all of them. In effect, this means the genome defines both a design variation and the arrangement for that design.

For example, consider the bold term within the BOP E-graph in Figure 5. The genome for that term is as follows, where * could be any integer since that e-class is not used by the term:

E_1, E_2	E_3	E_4	E_5	E_6	E_7	E_8	E_9
0	1	0	*	*	0	0	*

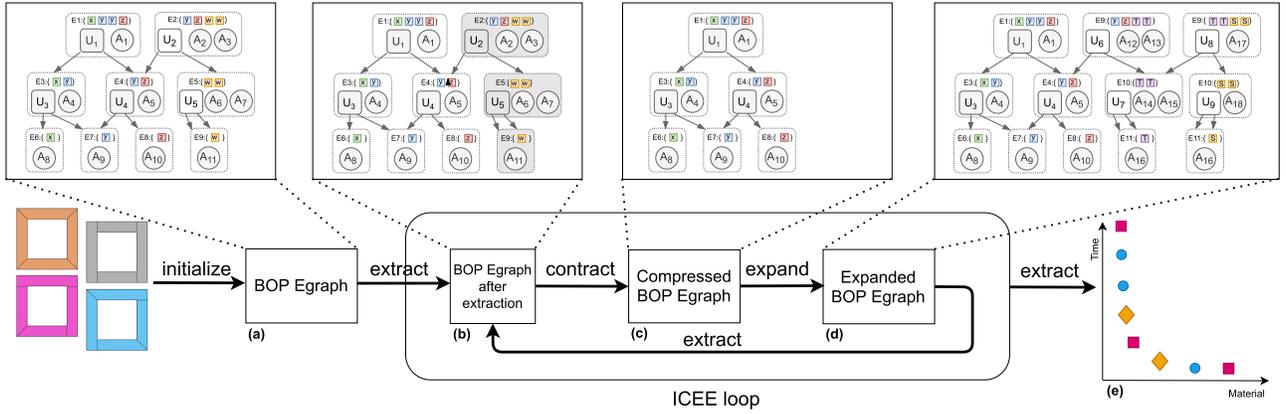


Fig. 6. Algorithm overview used the example in Figure 5. The first step initializes a BOP E-graph (Sections 4.3.2 and 4.3.3) with several design variants and a small number of fabrication arrangements (a). U and A represent union and atomic e-nodes respectively. As part of the ICEE loop, the algorithm extracts a Pareto Front (Section 4.3.4), which is used to score the e-classes in the BOP E-graph (b). For example, the gray e-class containing a “U” and an “A” e-node indicates a low score, i.e., the e-class did not contribute to Pareto-optimal solutions. The BOP E-graph is then contracted (Section 4.3.5) by removing the low-scored e-classes (and their parent e-nodes) to get a compressed BOP E-graph (c). As described in Section 4.3.6, this contracted BOP E-graph is then further expanded (d) by exploring more design variants and fabrication arrangements. The algorithm exits the loop when the termination conditions are reached, returning the final Pareto Front (e).

The root e-classes E_1 and E_2 share a single integer 0, meaning that the genome chooses the 0th e-node *across both* e-classes, and that it uses the first of the two design variants. Since this encoding boils down to a list of integers, which is valid as long as each integer corresponds to an e-node in that e-class, we can use simple mutation and single-point crossover operations.

A term does not completely determine a fabrication plan; it only specifies the arrangement. We need to additionally compute the cutting order for a given term to define a solution s and then evaluate the fabrication costs. We observe that the material cost does not depend on the cutting order and that precision and fabrication costs strongly correlate once the arrangement is fixed. This is not surprising since cutting orders that minimize set-ups will jointly reduce time and precision error. Given this observation, we can compute two solutions for each term, using two single-objective optimizations for computing cutting order: one that minimizes precision, and the other fabrication time.

We use two strategies to speed up these optimizations: (1) storing computed cutting orders in atomic e-nodes that will be shared across many terms and (2) a branch and bound technique. The optimization works as follows. Given a term, we first compute the optimal plans for the atomic e-nodes that have not been previously optimized. For each such e-node, we try to generate maximal P different orders of cuts, then extract the optimal plans with [Wu et al. 2019] method. We use this result to compute an upper and a lower bound for the term. If the lower bound is not dominated by the Pareto front of all computed solutions S , we run an optimization that uses the upper bound as a starting point (see Section 1.4 of the supplemental material for details).

We again terminate the algorithm if there is no hypervolume improvement within t_p iterations, or if we exceed mt_p iterations. In our experiments, we set $t_p = 20$ and $mt_p = 200$ and set the probability of crossover (mc_p) and mutation (mm_p) are set to be 0.95, 0.8 respectively.

4.3.5 BOP E-graph Contraction. As the algorithm proceeds, BOP E-graph contraction keeps the data structure from growing too large. To contract the BOP E-graph, we search for e-classes that represent bags of parts that have been sufficiently explored by the algorithm but are not present in Pareto-optimal designs. This indicates that we have already discovered the best way to fabricate these bags of parts but they still do not contribute to Pareto optimal solutions; these e-classes are then deleted.

To measure how much an e-class has been explored, we first compute how many variations of fabrication arrangements have been encoded in the BOP E-graph. This number is stored over the e-graph and updated after each expansion step to ensure consistency following contraction steps. The exploration score, E_{score} , is then defined as this value divided by the number of possible fabrication arrangements for an e-class, which we approximate by the number of parts in the e-class multiplied by the number of orientations of each part that can be assigned to the stock lumber.

The impact of an e-class, I_{score} , is measured based on how often it is used in the set of solutions in the current Pareto front. We use the assignment of solutions s to layers determined by the non-dominated sorting Section (3.1) to compute I_{score} for a given e-class. We initialize a I_{score} with value 0 and increment it by 10^{M-l} every time this e-class is used in a solution from layer l , where M is the total number of valid layers.

We normalize all computed exploration and impact scores to be between 0 and one and then assign the following pruning score to each e-class:

$$P_{score} = w \cdot I_{score} + (1 - w) \cdot (1 - E_{score}), w \in [0.0, 1.0],$$

where the weight w is chosen to trade-off between exploration and impact. If the P_{score} is smaller than the pruning rate, P_{rate} , the e-class is removed along with any e-nodes pointing to this e-class (i.e. parent e-nodes). We set w and P_{rate} to 0.7 and 0.3 in our implementation.

Table 1. Statistics for Each Input Model, Showing the Complexity in the Number of Parts (n_p), Number of Connectors (#C), Number of Connecting Variations (#CV), and Number of Design Variations That Define Unique BOP \mathcal{D}

Model	n_p	#C	#CV	$ \mathcal{D} $	Model	n_p	#C	#CV	$ \mathcal{D} $
Frame	4	4	22	13	A-Chair	18	3	6	4
L-Frame	6	8	16	65	F-Pot	8	1	4	4
A-Bookcase	12	6	16	192	Z-Table	15	6	16	63
S-Chair	14	14	32	66438	Loom	18	4	10	36
Table	12	10	24	1140	J-Gym	23	8	16	54
F-Cube	12	8	23	5	D-Chair	17	10	22	2280
Window	12	16	32	10463	Bookcase	15	22	44	65536
Bench	29	6	14	57	Dresser	10	10	25	480

4.3.6 BOP E-graph Expansion. We expand the BOP E-graph by first generating new design variations and then by generating fabrication arrangements for both the existing and newly generated design.

We generate new design variations using a single step of a genetic algorithm that operates over the design space. The probability of crossover (mc_d) and mutation (mm_d) are set to be 0.95, 0.8 respectively. We select the parent design variations from \mathcal{S} based on the non-dominated sorting technique (Section 3.1). Since many solutions in \mathcal{S} can correspond to the same design, we assign designs to the lowest layer that includes that design. We then generate new design variations with crossover and mutation operations. We use an integer vector encoding for each design. This time, the vector indexes the joint variations, e.g., for the designs shown in Figure 4, $d_1 = [0, 2, 1, 0]$, $d_2 = [1, 0, 0, 2]$. We get $K_m \cdot K_d$ design variations by applying K_m times of the single-step genetic algorithm. Then we apply the same heuristic done during initialization (Section 4.3.2), selecting the top K_{nd} , $K_{nd} \in [0, k_d]$. Finally, the resulting K_{nd} designs are included in the BOP E-graph. We set $K_m = 10$ in our implementation.

We generate fabrication arrangements for each of the new design variations using the algorithm described in Section 4.3.3, and they are added to the BOP E-graph maintaining the maximal sharing property. We further generate fabrication arrangements for existing design variations, using a similar scoring function used during contraction. This is done in two steps. First, we select root e-classes to expand based only on their impact score; namely, we take the top K_d root e-classes using non-dominated sorting. We then proceed to generate $K_f \times K_d$ fabrication arrangements using the algorithm described in Section 4.3.3). However, instead of generating the same number of fabrication arrangements variations for every selected root e-class, the number is adaptive to their pruning scores P_{score} (as defined in Section 4.3.5).

5 RESULTS AND DISCUSSION

In order to gauge the utility of our tool, we want to answer the following questions:

- (1) How much does searching the design space with the fabrication space improve generated fabrication plans?
- (2) How does our tool compare with domain experts who are asked to consider different design variations?
- (3) How does our tool's performance compare to a baseline naïve approach?

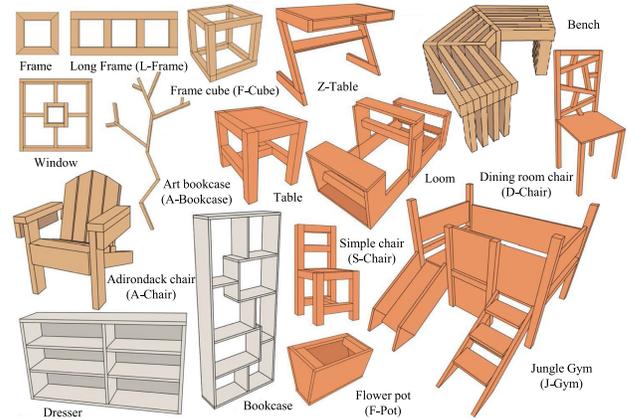


Fig. 7. Models used for all experiments in Section 5. Brown is used to indicate the models which are only made from 1D sequential cuts of lumber. Gray is for only from 2D partitioning of sheets. Orange is for both using 1D sequential cuts of lumber and 2D partitioning of sheets.

5.1 Models

We evaluate our method using the examples in Figure 7. Statistics for each model are shown in Table 1. These models vary widely in visual complexity and materials used—some are made from 1D sequential cuts on lumber, where others require 2D partitioning of sheets. Note the complexity of the search is not inherent to the visual complexity of the model, rather, it is determined by the number of connecting variations and the number of arrangements, which defines the size of the design space and the space of fabrication plans, respectively. For example, the Adirondack chair is more visually complex than the simple chair in Figure 7, but because it has about 5,000 times fewer design variations, it converges much more quickly. Models of Art bookcase, Dining room chair, F-Pot, Z-Table, Bench, and Adirondack chair are taken from [Wu et al. 2019].

5.2 Running environment

The parameters used in our ICEE algorithm are scaled based on the complexity of each model, measured in terms of the number of parts n_p and the size of the design space $|\mathcal{D}|$. We further introduce a single tuning parameter $\alpha \in [0.0, 1.0]$, which allows us to trade-off between exploring more design variations (smaller values of α) versus exploring more fabrication plans for given design variations (larger values). For all our experiments, we set α to the default value of 0.75. The ICEE parameters are set as follows: $K_d = 2^{\lceil \log_{10} |\mathcal{D}| \rceil}$, $N_{pop} = 4 \cdot K_d$, $K_f = \beta \cdot n_p$, $K_{nd} = \lfloor (1.0 - \alpha) \cdot K_d \rfloor$, and $P = 2 \cdot (\beta - 2)$, $t_d = 10$, $mt_d = 200$, $mc_d = 0.95$, $mm_d = 0.80$, $T = 4$ hours, $t_p = 20$, $mt_p = 200$, $mc_p = 0.95$, $mm_p = 0.80$, $w = 0.7$, $Prate = 0.3$ and $K_m = 10$, where $\beta = \lfloor 44 \cdot \alpha^7 + 2 \rfloor$. A table S1 of the supplemental material lists all parameters used in our algorithm.

We report the running times of our algorithm in Table 2 for the models in Figure 7. The above times are measured on a MAC laptop computer with 16 GB RAM and a 2.3 GHz 8-Core Intel Core i9 CPU. More discussion of the running time is in the supplemental material.

Table 2. Some Statistics and Running Times for Our ICEE Algorithm

Model	#O	#Iter	#EDV	#Arr	#PDV	CEt(m)	Et(m)	Total(m)
Frame	2	11	8	181	3	0.7	2.1	2.8
L-Frame	2	24	19	2,818	3	2.1	6.1	8.2
A-Bookcase	3	25	25	28,700	3	20.5	228.6	249.0
S-Chair	2	15	136	35,656	6	27.6	122.0	149.6
Table	2	18	50	9,346	9	5.9	34.9	40.8
F-Cube	2	23	4	3,499	3	1.4	4.0	5.5
Window	2	23	116	81,026	4	32.8	98.9	131.7
Bench	2	25	16	37,436	3	30.3	215.1	245.4
A-Chair	2	28	4	14,440	3	3.1	9.6	12.7
F-Pot	3	14	3	185	2	1.7	13.0	14.7
Z-Table	3	70	41	336,091	6	17.1	71.1	88.2
Loom	3	21	10	1,812	5	3.1	74.6	77.7
J-Gym	3	46	18	286,239	3	37.0	72.0	109.0
D-Chair	2	18	40	15,054	7	27.7	228.8	256.5
Bookcase	3	15	32	34,756	11	39.4	336.8	376.3
Dresser	3	20	44	22,209	5	14.1	241.2	255.4

For each model, we first report the number of targeting objectives (#O) where 2 indicates material usage (f_c) and fabrication time (f_t), and 3 indicates all of the three objectives including cutting precision (f_p). We also report the number of iterations (#Iter), explored design variations (#EDV) and arrangements (#Arr), and Pareto front design variations (#PDV). We report the running time of BOP E-graph contraction and expansion (CEt), and Pareto front extraction (Et), as well as the total time. All running times are in minutes.

5.3 Benefits of Design Exploration

To demonstrate the benefit of simultaneous exploration of the design variation and fabrication plan spaces, we compare our tool against optimizing the fabrication plan for a single design.

Figure 8 shows the comparison between our pipeline and the Carpentry Compiler pipeline [Wu et al. 2019], which only considers a single design. The parameter setting of their pipeline and additional results can be found in Section 2 of the supplemental material. We explore the trade-offs for fabrication time and material usage for the designs where all cuts can be executed with standard setups (these are considered to have no precision error) and include a third objective of precision error for the models where that is not possible. The Pareto fronts across designs generated by our tool cover a larger space and many of our solutions dominate those from previous work.

Exploring design variations enables better coverage of the Pareto front, which enables finding better trade-offs. These trade-offs are lower-cost overall, cover more of the extrema, and are more densely available. For example, a hobbyist may want to minimize material cost independent of time, as the fabrication process is enjoyable, and they consider it to have no cost. Material cost is hard to save, but our exploration of design variations enable solutions that reduce material cost by 7% in the Loom, 7% in the Jungle Gym, 15% in the Frame, and 25% in the Bookcase. On the other hand, someone with free access to reclaimed wood may only care about the total fabrication time. Our approach enables solutions that reduce fabrication time by 60%—two models saved between 50%–60%, three between 30%–35%, and four between 20%–30%, for example— a huge time-saving. If creating a very precise model is imperative, and a user would take great care to manufacture it exactly, then for four models, we find solutions that reduce error by 61%–77%. The detailed data are listed in Table S6 of the supplemental material.

Some examples don't lie at the extrema: businesses often need to find a balance between the cost of materials, time spent on a project, and overall project quality, and the particular trade-off will depend on their accounting needs. Our method enables finding solutions with better trade-offs. Concretely, consider a carpenter charging \$40/h. When scalarizing our multi-objective function into this single objective of money, we have several examples

where the lowest cost on our Pareto front is 5%–8% cheaper than the lowest cost on the baseline Pareto front, such as the Z-Table, Flowerpot, Jungle Gym, Dresser, Bookcase, and Art Bookcase. The window and frame have cost savings of 12% and 20%, respectively. Though a cost reduction of several percent might appear insignificant, in production at scale, it represents thousands of dollars potentially saved. This scalarization function is just one way for a user to judge the trade-off between different aspects of the multi-objective function. In reality, the user probably has some notion of what trade-off would be ideal for their purposes, and will use the Pareto front to represent the full space of options and make an informed choice. This scalarized tradeoff is further examined in Table S8 of the supplemental material.

Figure 9 highlights how exploring design variations generates fabrication plans that can dominate those generated from no design variation exploration. Figure 10 then demonstrates how design variations enable diverse trade-offs that save on different costs.

5.4 Comparison with Experts

For each model, we asked carpentry experts to generate design variations and fabrication plans. The resulting points are plotted as diamonds in Figure 8. Since experts produce each solution by hand, they produced Pareto fronts with many fewer solutions than our tool. For 14 of 16 models (except the Loom and Dresser models), solutions generated by our tool dominate the expert solutions. This suggests that, generally, although expert plans seem sensible, our tool generates better output thanks to its ability to generate more design variations and fabrication plans, including potentially unintuitive packing or cutting orders, and evaluate them much more quickly than a human.

5.5 Performance Evaluation

To test whether the BOP E-graph's sharing is important for our tool's performance, we compare against a nested-optimization pipeline built on the Carpentry Compiler [Wu et al. 2019]. The baseline approach invokes the Carpentry Compiler pipeline on each design variant that our tool explores, and then it takes the Pareto front across all design variations.

We choose five models of varying complexity to evaluate performance and show result in Table 3. We tuned the parameters of the baseline method so we could achieve results that were as close as possible, if not qualitatively the same (when the baseline method ran to completion). Full results are available in the supplemental material (Table S7 and Figure S1). This indicates that our co-optimization approach yields similar results to the nested approach over the same space. When it comes to performance, our approach is about one order of magnitude faster. We attribute this speedup to the sharing captured by the BOP E-graph; we only had to evaluate common sub-designs and sub-fabrication-plans one time, whereas the baseline shared no work across its different invocations for each design variant.

5.6 Fabricated Results

We validated our pipeline by physically constructing some of the models according to the design variation-fabrication plan pairs generated by our tool. Figure 11 shows the results.

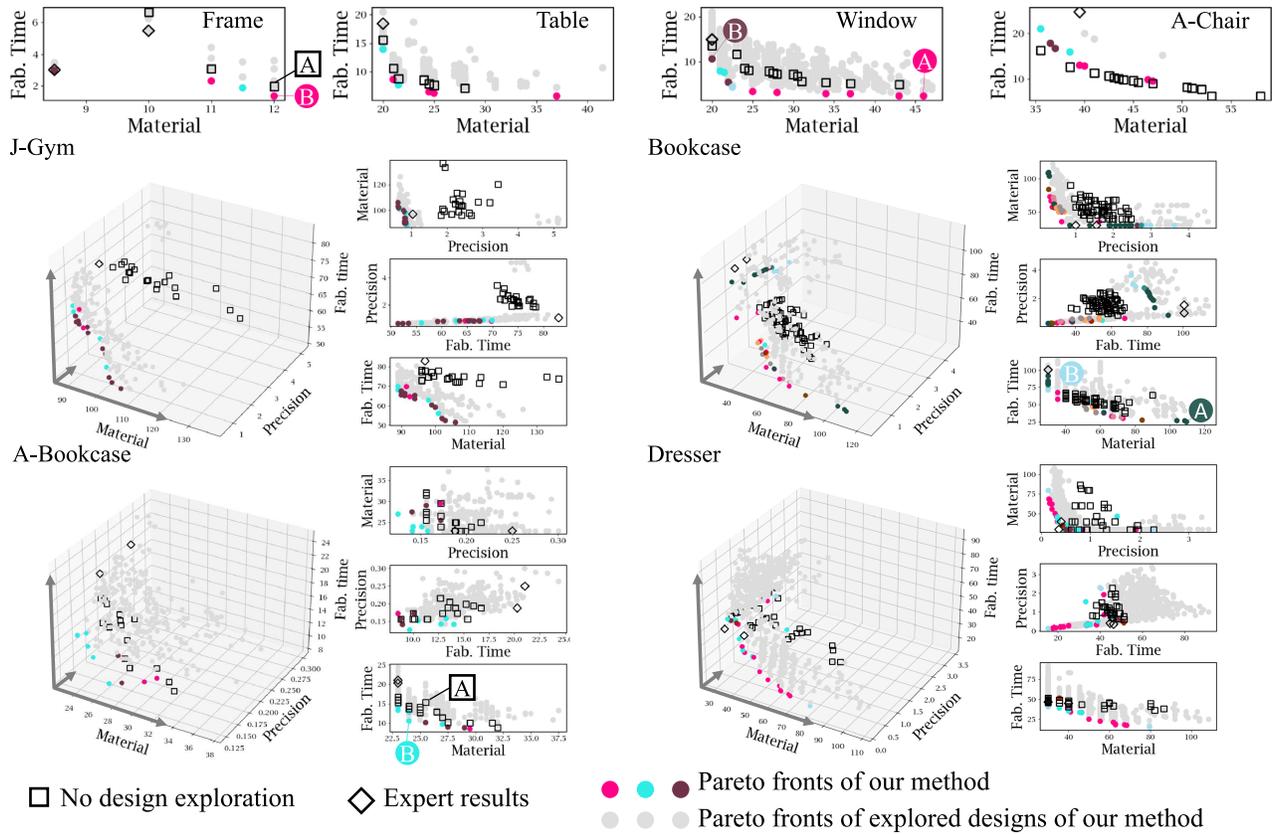


Fig. 8. Pareto fronts are computed from our pipeline with design optimization as colored dots. Each color corresponds to a different design. The gray dots indicate the Pareto fronts of all explored design variations. These are compared against Pareto fronts computed without design optimization (fabrication optimization only, using the original model as the input design) as squares, and expert fabrication plans as diamonds. Often, fabrication plans from a design variant are more optimal than those generated from an input design. For the unit of objective metrics, material usage (f_c) is in dollars, cutting precision (f_p) is in inches, and fabrication time (f_t) is in minutes. Some (design, fabrication plan) pairs indicated with capital letters are visualized in Figures 9 and 10.

6 DISCUSSION

6.1 Multi-Materials and Cutting Tools

Mechanical or aesthetic reasons might motivate designers to combine multiple materials, such as different types of wood, or wood and metal, in one model. Adding new materials to our approach involves almost no implementation overhead: we must select which cutting tools are appropriate, and accommodate the material's costs into our metrics. Then, we simply need to indicate which material a given part is made of, exactly the same way we designate whether parts belong on 1D lumber or 2D stock. As shown in Figure 12, we have created a mixed-material model to showcase our ability to handle this added design complexity. The loom is made of two different types of wood as well as one kind of metal. All parts are optimized in the same e-graph and treated identically to the base implementation. We describe the cost metrics for different materials in the supplemental material (Section 1.3.1).

6.2 Objectives

Our method also naturally extends to other objective functions. We show one example in Figure 13, where we consider stability as an

Table 3. Results of the Performance Validation Experiment

Model	$ \mathcal{D} $	#EDV	Time (min)	
			Ours	Baseline
Frame	13	8	2.8	6.5
Jungle Gym	54	18	109.0	761.2
Long frame	65	19	8.2	59.7
Table	1140	59	40.8	612.8
Window	10463	116	131.7	2050.0

"Ours" indicates the ICEE algorithm of this article. "Baseline" indicates extracting the Pareto front fabrication plans for each design variation explored by our method independently with the Carpentry Compiler pipeline [Wu et al. 2019]. The size of design space $|\mathcal{D}|$ and the number of explored design variations (EDV) are also reported. Our method and the baseline method produce two Pareto fronts which are indistinguishable. This conclusion is not shown here; direct comparisons of hypervolume can be non-intuitive due to the scale and how hypervolume is measured. Please refer to the supplemental material (Figure S1), which contains plots comparing the results of the two methods. Even with identical results, our time improvement is significant.

additional objective which we calculate with physical simulation. Notably, stability is invariant to the fabrication plan, and depends solely on the design itself, so it only needs to be measured once, at the root node. However, two designs can have different stability costs but share the same BOP. Figure 13 (a) and (b), exhibits one BOP which captures two different designs.

In this example, since the other metrics (time and material cost) do not exhibit this dependency, we can simply assign to the

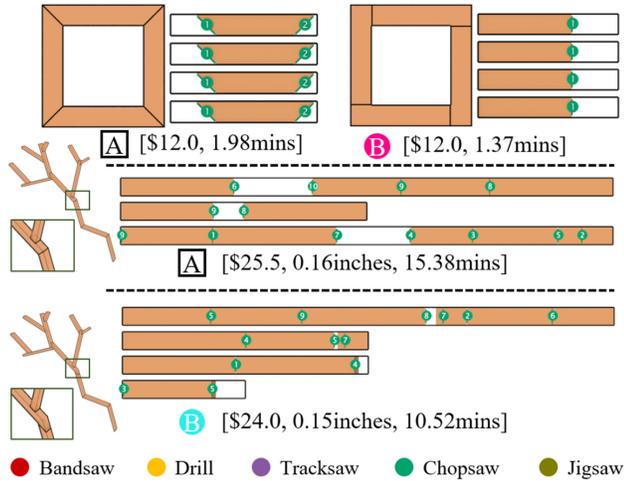


Fig. 9. Two examples where searching the design space revealed fabrication plans that completely dominated the fabrication plans generated for the input design. With the design variations, our pipeline could search for a design variation of the frame, which turns all angled cutting to vertical. With Design B, we find a fabrication plan which takes less time than the least time-consuming plan A of the input design. Similarly, we show two fabrication plans of the A-Bookcase model where the design and fabrication plan B dominate the input design A. The fabrication costs are indicated in the figure with the order of material cost, precision error, and fabrication time. The cutting orders are labeled with colored dots and numbers, where colors indicate selected cutting tools, and stacked cuts are labeled with the same number.

root nodes the stability cost of the best-performing design that corresponds to that BOP; thus, the cost for any given BOP is the best cost of any design that is represented by that BOP. Note that fabrication plans depend solely on the BOP. In general, if we want to use more than one metric like this one—a metric that depends on the design, and is not completely determined by a term in the e-graph—we would need to compute the different trade-offs for the variations during extraction, as was done with cutting order and precision, described in Section 4.3.4.

6.3 Convergence

While our results show the significance of the approach to reduce fabrication cost in practice, we cannot make any guarantees that the plans we output are on the globally-optimal Pareto front. Indeed, we do not anticipate that any alternative approach would be able to have such strong guarantees given the inherent complexity of the problem. This convergence limitation impacts our method in three different ways.

Parameter Tuning. Due to limitations in exploring the full combinatorial space, parameters of our search algorithm may influence convergence. Because the key aspect of ICEE is simultaneously searching “broad” (design variations) and “deep” (fabrication plans for various designs), we expose the α parameter that trades-off between depth and breadth during search. Exposing this single parameter, enables us to improve performance in special circumstances. For example, when not much can be gained from design variations, a larger α will enable searching deeply on a particular

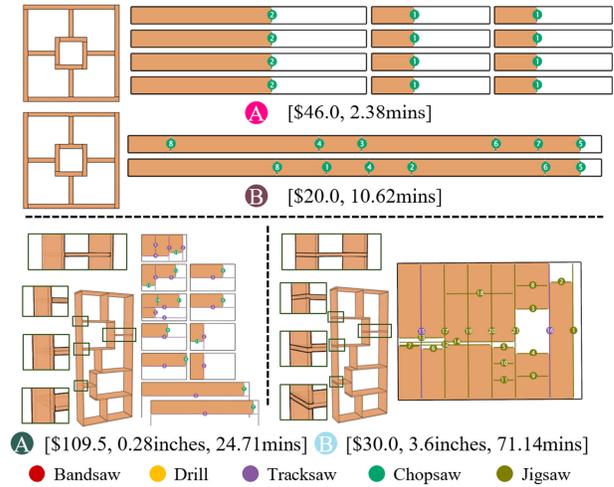


Fig. 10. Two examples where exploring different designs lead to a wider diversity of plans, where each trade-off on the Pareto front is only possible because of the underlying design. The window provides a simpler example. Design A is very uniform, with only three distinct parts. This design makes it easy to save on fabrication time because we can stack the cuts across different stocks. Design B features more varied cuts, unlike A, where each of the sides was the same length. This irregularity allows all the parts to be effectively rearranged onto just two pieces of stock. Regular pieces would not fit as nicely and result in wastage. Material cost is very low, but because of the tight packing, much more time is needed to make each individual cut. The bookcase example showcases how some unintuitive design decisions lead to cost savings. In this example, Design A’s two long, identical side pieces mean more opportunities for stacking, of which the fabrication plan takes full advantage. This design enables a very low time cost, but uses a lot of material. Design B’s left side is broken up by the shelves, and without a second long piece, it is possible to pack all the pieces onto a single piece of lumber. Here, the material used is economical, but the carpenter must take time to cut pieces from a complex layout.



Fig. 11. Fabrication results of two window variations. The different designs and fabrication plans trade-off fabrication time and material usage.

design finding better solutions. All the results shown in this article use the default value for α that we have found effective in practice.

Comparison with Wu et al. [2019]. The fundamental difference between our work and [Wu et al. 2019] is that incorporating more design variations increases the design space, enabling us to find better performing results. Since the search space of this prior work is a subset of the search space we explore, our results should be strictly better. However, since neither method can ensure the

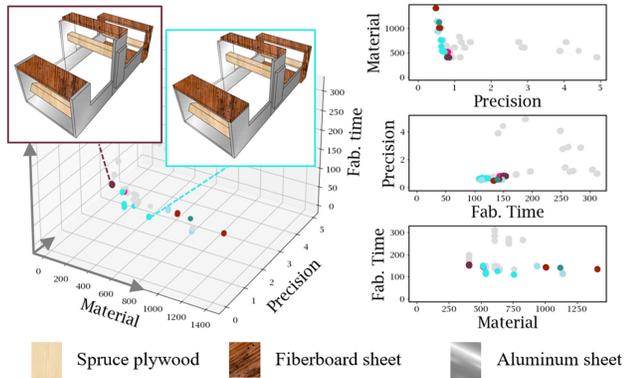


Fig. 12. A loom model with mixed material where two kinds of wood (spruce plywood and medium density fiberboard sheet) and one kind of metal (aluminum sheet) are assigned to each part.

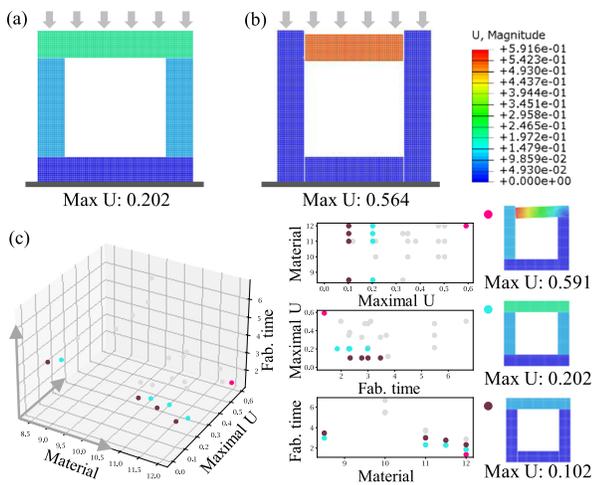


Fig. 13. Pareto fronts computed from by our pipeline for the Frame model with three objective functions, material usage f_c , fabrication time f_t and stability performance. The physical stability of each design variation is simulated with Abaqus/CAE 2021, measured with the maximal displacement (Max U). All displacements are in inches. In this figure, (a) is the displacement visualization in a direction, (b) is the displacement visualization of the same design but with a different direction, and (c) plots the Pareto fronts computed from our pipeline where three design variations are selected.

results lie on the true Pareto front due to limitations in convergence, tuning parameters of both approaches may influence this result. An example of this limitation is shown in the A-Chair example in Figure 7. We show in the supplemental material (Section 2.4) how tuning α to explore more deeply improves this result and also report experiments for tuning the four parameters from [Wu et al. 2019].

Increasing the Design Space. A final implication of the intractable search is that it is possible to achieve worse results by increasing the design space in special circumstances. We discuss in the supplemental material (Section 2.4) an example where we make the design space 145 times larger by including variations that do not benefit the search.

6.4 Limitations and Future Work

Our current approach encodes only discrete design variants in the BOP E-graph. An interesting direction for future work would be to support continuous variations in the designs space, which can provide a larger space of fabrication plans to explore. Such variations could result in designs that do not preserve the original shape but have similar appearance and functionality. However, supporting continuous design variants in an e-graph would require designing a new metric for comparing two design variants for equivalence. This is challenging because e-graphs heavily exploit transitivity, so any error in the metric could lead to arbitrarily different designs being considered “equivalent”.

Several steps of our algorithm can also be parallelized for performance (e.g., generating design variants)—we leave this as an extension for the future.

Another direction we are eager to explore is accounting for other factors in the Pareto front. Currently, our technique finds a design variant and fabrication plan that optimizes fabrication time, material cost, and precision. Other interesting factors that can guide the search include stability measurement of different design variations and assembly considerations. The assembly process will not only impact the total fabrication time and cost but should be considered when computing the structural soundness and durability of the final product. We are also interested in exploring more complex wood connectors (integral joints).

We are also keen to explore broader applications of the ICEE strategy for integrating feedback-directed search in other e-graph-based optimization techniques including additive and subtractive manufacturing. Past work applying e-graphs for design optimization in CAD [Nandi et al. 2020] and for improving accuracy in floating-point code [Panchevka et al. 2015] have relied on ad hoc techniques for controlling the growth of the e-graph, e.g., by carefully tuning rules used during rewrite-based optimization. We hope to explore whether ICEE can help with optimization in such domains by focusing the search on more-profitable candidates and easing implementation effort by reducing the need for experts to carefully tune rewrite rules.

The most time-consuming part of our ICEE algorithm lies in the Pareto front extraction phase. A pruning strategy with learning-based methods for predicting the objective metrics of an arrangement might be an interesting and valuable area of research.

7 CONCLUSION

We have presented a new approach to co-optimizing model design variations and their fabrication plans. Our approach relies on the insight that fabrication plans across design variants will share similar structure. We capture this sharing with the BOP E-graph data structure that considers fabrication plans equivalent if they produce the same BOP. The BOP E-graph also lets us guide the search toward profitable design variants/fabrication plans with a technique we call ICEE that may be useful for the uses of e-graphs in other applications. Results generated by our tool compare favorably against both expert-generated designs and a baseline built using prior work, indicating that the sharing captured by the BOP E-graph is essential to efficiently exploring the large, combined space of design variants and fabrication plans.

ACKNOWLEDGMENTS

The authors would like to thank anonymous reviewers for their helpful feedback; Haomiao Wu for her contribution to the algorithm development in the early stage of the project; Elias Baldwin, David Tsay, Alexander Lefort, and Qiyang Tan for helping the experiments.

REFERENCES

- Thomas Alderighi, Luigi Malomo, Daniela Giorgi, Bernd Bickel, Paolo Cignoni, and Nico Pietroni. 2019. Volume-aware design of composite molds. *ACM Transactions on Graphics* 38, 4 (2019), 1–12.
- Anne Auger, Johannes Bader, Dimo Brockhoff, and Eckart Zitzler. 2009. Theory of the hypervolume indicator: Optimal μ -distributions and the choice of the reference point. In *Proceedings of the 10th ACM SIGEVO Workshop on Foundations of Genetic Algorithms*. 87–102.
- Bernd Bickel, Paolo Cignoni, Luigi Malomo, and Nico Pietroni. 2018. State of the art on stylized fabrication. *Computer Graphics Forum* 37, 6 (2018), 325–342. DOI: <https://doi.org/10.1111/cgf.13327>
- Paolo Cignoni, Nico Pietroni, Luigi Malomo, and Roberto Scopigno. 2014. Field-aligned mesh joinery. *ACM Transactions on Graphics* 33, 1 (2014), 1–12.
- Kalyanmoy Deb. 2014. Multi-objective optimization. *Search Methodologies*. Springer.
- Kalyanmoy Deb and Himanshu Jain. 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation* 18, 4 (2013), 577–601.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- Kalyanmoy Deb and Ankur Sinha. 2009. Solving bilevel multi-objective optimization problems using evolutionary algorithms. In *Proceedings of the International Conference on Evolutionary Multi-criterion Optimization*. Springer, 110–124.
- Stephan Dempe. 2018. *Bilevel Optimization: Theory, Algorithms and Applications*. TU Bergakademie Freiberg, Fakultät für Mathematik und Informatik.
- Simon Duenser, Roi Poranne, Bernhard Thomaszewski, and Stelian Coros. 2020. RoboCut: Hot-wire cutting with robot-controlled flexible rods. *ACM Transactions on Graphics* 39, 4 (2020), 98–1.
- Gabriele Eichfelder. 2010. Multiobjective bilevel optimization. *Mathematical Programming* 123, 2 (2010), 419–449.
- Jimmy Etienne, Nicolas Ray, Daniele Panozzo, Samuel Hornus, Charlie CL Wang, Jonás Martínez, Sara McMains, Marc Alexa, Brian Wyvill, and Sylvain Lefebvre. 2019. CurviSlicer: Slightly curved slicing for 3-axis printers. *ACM Transactions on Graphics* 38, 4 (2019), 1–11.
- Chi-Wing Fu, Peng Song, Xiaoqi Yan, Lee Wei Yang, Pradeep Kumar Jayaraman, and Daniel Cohen-Or. 2015. Computational interlocking furniture assembly. *ACM Transactions on Graphics* 34, 4 (2015), 11 pages. DOI: <https://doi.org/10.1145/2766892>
- Akash Garg, Alec Jacobson, and Eitan Grinspun. 2016. Computational design of reconfigurables. *ACM Transactions on Graphics* 35, 4 (2016), 90–1.
- Konstantinos Gavriil, Ruslan Guseinov, Jesús Pérez, Davide Pellis, Paul Henderson, Florian Rist, Helmut Pottmann, and Bernd Bickel. 2020. Computational design of cold bent glass façades. *ACM Transactions on Graphics* 39, 6 (2020), 1–16.
- Werner Halter and Sanaz Mostaghim. 2006. Bilevel optimization of multi-component chemical systems using particle swarm optimization. In *Proceedings of the 2006 IEEE International Conference on Evolutionary Computation*. IEEE, 1240–1247.
- Kristian Hildebrand, Bernd Bickel, and Marc Alexa. 2013. Orthogonal slicing for additive manufacturing. *Computers & Graphics* 37, 6 (2013), 669–675.
- Rajeev Joshi, Greg Nelson, and Keith Randall. 2002. Denali: A goal-directed superoptimizer. *SIGPLAN Notices* 37, 5 (2002), 304–314. DOI: <https://doi.org/10.1145/543552.512566>
- Bongjin Koo, Jean Hergel, Sylvain Lefebvre, and Niloy J. Mitra. 2017. Towards zero-waste furniture design. *IEEE Transactions on Visualization and Computer Graphics* 23, 12 (2017), 2627–2640. DOI: <https://doi.org/10.1109/TVCG.2016.2633519>
- Bongjin Koo, Wilmot Li, JiaXian Yao, Maneesh Agrawala, and Niloy J. Mitra. 2014. Creating works-like prototypes of mechanical objects. *ACM Transactions on Graphics* 33, 6 (2014), 1–9.
- Manfred Lau, Akira Ohgawara, Jun Mitani, and Takeo Igarashi. 2011. Converting 3D furniture models to fabricatable parts and connectors. In *Proceedings of the ACM SIGGRAPH 2011 Papers*. ACM, New York, NY, 6 pages. DOI: <https://doi.org/10.1145/1964921.1964980>
- Danny Leen, Tom Veuskens, Kris Luyten, and Raf Ramakers. 2019. JigFab: Computational fabrication of constraints to facilitate woodworking with power tools. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 12 pages. DOI: <https://doi.org/10.1145/3290605.3300386>
- Jie Lu, Jialin Han, Yaoguang Hu, and Guangquan Zhang. 2016. Multilevel decision-making: A survey. *Information Sciences* 346, C (2016), 463–487. <https://dl.acm.org/doi/abs/10.1016/j.ins.2016.01.084>
- Ali Mahdavi-Amiri, Fenggen Yu, Haisen Zhao, Adriana Schulz, and Hao Zhang. 2020. VDAC: Volume decompose-and-carve for subtractive manufacturing. *ACM Transactions on Graphics* 39, 6 (2020), 1–15.
- Henrique Teles Maia, Dingzeyu Li, Yuan Yang, and Changxi Zheng. 2019. LayerCode: Optical barcodes for 3D printed shapes. *ACM Transactions on Graphics* 38, 4 (2019), 1–14.
- Chandrakana Nandi, Max Willsey, Adam Anderson, James R. Wilcox, Eva Darulova, Dan Grossman, and Zachary Tatlock. 2020. Synthesizing structured CAD models with equality saturation and inverse transformations. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation*. Alastair F. Donaldson and Emina Torlak (Eds.), ACM, 31–44. DOI: <https://doi.org/10.1145/3385412.3386012>
- Charles Gregory Nelson. 1980. *Techniques for Program Verification*. Ph.D. Dissertation. Stanford University, Stanford, CA, AAI8011683.
- Pavel Panekha, Alex Sanchez-Stern, James R. Wilcox, and Zachary Tatlock. 2015. Automatically improving accuracy for floating point expressions. *SIGPLAN Notices* 50, 6 (2015), 1–11. DOI: <https://doi.org/10.1145/2813885.2737959>
- Varot Premtoon, James Koppel, and Armando Solar-Lezama. 2020. Semantic code search via equational reasoning. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. Association for Computing Machinery, New York, NY, 1066–1082. DOI: <https://doi.org/10.1145/3385412.3386001>
- Yuliy Schwartzburg and Mark Pauly. 2013. Fabrication-aware design with intersecting planar pieces. In *Proceedings of the Computer Graphics Forum*. Wiley Online Library, 317–326.
- Xinping Shi and Hong Sheng Xia. 2001. Model and interactive algorithm of bi-level multi-objective decision-making with multiple interconnected decision makers. *Journal of Multi-Criteria Decision Analysis* 10, 1 (2001), 27–34.
- Ankur Sinha, Pekka Malo, and Kalyanmoy Deb. 2017. A review on bilevel optimization: From classical to evolutionary approaches and applications. *IEEE Transactions on Evolutionary Computation* 22, 2 (2017), 276–295.
- Peng Song, Chi-Wing Fu, Yueming Jin, Hongfei Xu, Ligang Liu, Pheng-Ann Heng, and Daniel Cohen-Or. 2017. Reconfigurable interlocking furniture. *ACM Transactions on Graphics* 36, 6 (2017), 14 pages. DOI: <https://doi.org/10.1145/3130800.3130803>
- Michael Stepp, Ross Tate, and Sorin Lerner. 2011. Equality-based translation validator for LLVM. In *Proceedings of the 23rd International Conference on Computer Aided Verification*. Springer-Verlag, Berlin, 737–742.
- Pengbin Tang, Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. 2020. A harmonic balance approach for designing compliant mechanical systems with nonlinear periodic motions. *ACM Transactions on Graphics* 39, 6 (2020), 1–14.
- Ross Tate, Michael Stepp, Zachary Tatlock, and Sorin Lerner. 2009. Equality saturation: A new approach to optimization. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, 264–276. DOI: <https://doi.org/10.1145/1480881.1480915>
- Nobuyuki Umetani, Takeo Igarashi, and Niloy J. Mitra. 2012. Guided exploration of physically valid shapes for furniture design. *ACM Transactions on Graphics* 31, 4 (2012), 11 pages. DOI: <https://doi.org/10.1145/2185520.2185582>
- Yisu Remy Wang, Shana Hutchison, Jonathan Leang, Bill Howe, and Dan Suciu. 2020. SPORES: Sum-product optimization via relational equality saturation for large scale linear algebra. In *Proceedings of the VLDB Endowment*. 1919–1932. DOI: <https://doi.org/10.14778/3407790.3407799>
- Ziqi Wang, Peng Song, Florin Isvoranu, and Mark Pauly. 2019. Design and structural optimization of topological interlocking assemblies. *ACM Transactions on Graphics* 38, 6 (2019), 1–13.
- Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panekha. 2021. Egg: Fast and extensible equality saturation. In *Proceedings of the ACM on Programming Languages*. 29 pages. DOI: <https://doi.org/10.1145/3434304>
- Chenming Wu, Haisen Zhao, Chandrakana Nandi, Jeffrey I. Lipton, Zachary Tatlock, and Adriana Schulz. 2019. Carpentry compiler. *ACM Transactions on Graphics* 38, 6 (2019), 1–14.
- Jinfan Yang, Chrystiano Araujo, Nicholas Vining, Zachary Ferguson, Enrique Rosales, Daniele Panozzo, Sylvain Lefebvre, Paolo Cignoni, and Alla Sheffer. 2020. DHFSlicer: Double height-field slicing for milling fixed-height materials. *ACM Transactions on Graphics* 39, 6 (2020), 1–17.
- Yong-Liang Yang, Jun Wang, and Niloy J. Mitra. 2015. Reforming shapes for material-aware fabrication. In *Proceedings of the Computer Graphics Forum*. Wiley Online Library, 53–64.
- Yafeng Yin. 2000. Genetic-algorithms-based approach for bilevel programming models. *Journal of Transportation Engineering* 126, 2 (2000), 115–120.
- Xiaoting Zhang, Guoxin Fang, Mélina Skouras, Gwenda Gieseler, Charlie Wang, and Emily Whiting. 2019. Computational design of fabric formwork. *ACM Transactions on Graphics* 38, 4 (2019), 1–13.
- Haisen Zhao, Fanglin Gu, Qi-Xing Huang, Jorge Garcia, Yong Chen, Changhe Tu, Bedrich Benes, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. 2016. Connected fermat spirals for layered fabrication. *ACM Transactions on Graphics* 35, 4 (2016), 1–10.

Received July 2021; revised November 2021; accepted December 2021