# Grassland: A Rapid Algebraic Modeling System for Million-variable Optimization 

Xihan Li<br>University College London<br>London, The United Kingdom xihan.li@cs.ucl.ac.uk

Xiongwei Han<br>Huawei Noah's Ark Lab<br>Shenzhen, China<br>hanxiongwei@huawei.com

Zhishuo Zhou<br>Fudan University<br>Shanghai, China<br>zhouzs18@fudan.edu.cn

Mingxuan Yuan<br>Huawei Noah's Ark Lab<br>Shenzhen, China<br>yuan.mingxuan@huawei.com

Jia Zeng<br>Huawei Noah's Ark Lab<br>Shenzhen, China<br>zeng.jia@huawei.com

Jun Wang<br>University College London<br>London, The United Kingdom<br>jun.wang@cs.ucl.ac.uk


#### Abstract

An algebraic modeling system (AMS) is a type of mathematical software for optimization problems, which allows users to define symbolic mathematical models in a specific language, instantiate them with given source of data, and solve them with the aid of external solver engines. With the bursting scale of business models and increasing need for timeliness, traditional AMSs are not sufficient to meet the following industry needs: 1) million-variable models need to be instantiated from raw data very efficiently; 2) Strictly feasible solution of million-variable models need to be delivered in a rapid manner to make up-to-date decisions against highly dynamic environments. Grassland is a rapid AMS that provides an end-toend solution to tackle these emerged new challenges. It integrates a parallelized instantiation scheme for large-scale linear constraints, and a sequential decomposition method that accelerates model solving exponentially with an acceptable loss of optimality. Extensive benchmarks on both classical models and real enterprise scenario demonstrate $6 \sim 10 \mathrm{x}$ speedup of Grassland over state-of-the-art solutions on model instantiation. Our proposed system has been deployed in the large-scale real production planning scenario of Huawei. With the aid of our decomposition method, Grassland successfully accelerated Huawei's million-variable production planning simulation pipeline from hours to $3 \sim 5$ minutes, supporting near-real-time production plan decision making against highly dynamic supply-demand environment.


## CCS CONCEPTS

- Mathematics of computing $\rightarrow$ Mathematical software; • Applied computing $\rightarrow$ Operations research; Supply chain management.

[^0]
## KEYWORDS

algebraic modeling system, large-scale optimization

## ACM Reference Format:

Xihan Li, Xiongwei Han, Zhishuo Zhou, Mingxuan Yuan, Jia Zeng, and Jun Wang. 2021. Grassland: A Rapid Algebraic Modeling System for Millionvariable Optimization. In Proceedings of the 30th ACM Int'l Conf. on Information and Knowledge Management (CIKM '21), November 1-5, 2021, Virtual Event, Australia. ACM, New York, NY, USA, 12 pages. https://doi.org/10. 1145/XXXXXX.XXXXXX

## 1 INTRODUCTION

Mathematical optimization is a powerful analytics technology that allows companies to make optimal decisions based on available business data, widely applied in industry scenarios including logistics[9], manufacturing[6], finance[7] and energy. It abstracts key features of a complex business problem as an optimization model, which consists of objectives (business goal), variables (decisions to be made) and constraints (business rules). In such a way the business problem is decomposed into two stages: conversion - converting the problem to a canonical optimization model, and solving - finding the optimal or approximated solution of the model.

In practice, the conversion stage consists of two steps. The first step is modeling, which means writing down the expression of objective and constraints in a formulated way. E.g., using mathematical expression $\sum_{i \in S} x_{i} \leq c a p$ to formulate a capacity constraint that the total production amount of specific products should not exceed the plant capacity. The second step is instantiation, which means to generate a particular instance of the model when real data are available. E.g., when we know today's plant capacity is $c a p=100$ and available set of products is $S=\{1,3,5\}$, we get a particular instance of expression $x_{1}+x_{3}+x_{5} \leq 100$ for today (and tomorrow's instance might be very different). For the solving stage, we usually use solver engines that are sophisticatedly developed to solve specific kinds of models, such as Gurobi, CPLEX, Mosek and CLP. Decomposition methods may also apply when the model is large. A toy example of mathematical optimization pipelines for practical online business scenarios is shown in Figure 1, in which the high-level system to handle both the conversion and solving stages is usually called algebraic modeling system (AMS).
In recent years, the scale and complexity of business problems are dramatically increased, e.g., a large electronics company can involve over $10^{5}$ types of products and $10^{2}$ plants worldwide with


Figure 1: A toy example of mathematical optimization pipeline for practical business decision making scenarios.
different standard. As a result, their corresponding optimization models become extremely massive and cumbersome. They not only involve millions of decision variables, but also contain extremely lengthy real-world business constraints, which cover numerous case of real business logic (e.g., production hierarchy, inventory control, supply-demand modeling, delay minimization, intra/interfactory transshipment and replacement) thus can take hundreds of pages to document. While million-variable models can be a burden for solving stage, instantiating hundred-page business constraints efficiently as standard models is also highly nontrivial for conversion stage. For end-to-end optimization, both the two stages can be extremely time-consuming with traditional toolchain.

However, the information age calls for rapid optimization. Only in this way can business decisions be frequently adjusted and updated to reflect the latest state of fast-changing market, and fulfill customers' increasing need for timeliness. The more rapidly getting optimized decisions from latest available data, the more timely the company can respond to market change and other uncertainty factors. This especially applies to rolling horizon that a time-dependent model is solved repeatedly. E.g., a manufacturing company that can update its production plan in minutes (high re-planning periodicity) to handle unexpected urgent orders or factory stoppage is more competitive than its counterparts who have to wait for hours (low re-planning periodicity) for a new optimized plan. The same goes for other scenarios like logistics and finance, in which the timeliness of business decisions is directly related to user experience and profits, acting as a core competitiveness of modern enterprises.

Moreover, rapid optimization creates remarkable possibilities for business intelligence. First, it can act as an analytical tool that provides metrics for high-level business decisions. E.g., buyers can evaluate different raw material purchase plans via running a "simulated" production planning model for each plan, and check their corresponding order fulfillment rates. Second, it can serve as a basis for larger or more complicated optimization. E.g., to strictly conserve order priority constraint (which is nonlinear), we can run production planning multiple times, where high-priority orders are planned before low-priority ones. All these build on the cornerstone that a single shot of end-to-end optimization can be very rapid.

Sadly, current state-of-the-art AMSs are far from supporting the aforementioned ambition of rapid optimization. In Figure 1, for the conversion stage, they lack a principled design to stress
the efficiency of model instantiation, especially ignoring the parallelization and vectorization of operations. This issue is minor for normal-sized, simple models, but significantly raised as a major bottleneck for million-variable, hundred-page documented business scenarios. For the solving stage, they focus on lossless decomposition methods such as Benders decomposition, which are not widely applicable since the special block structures they required are easily broken in complex real scenarios. As a result, both two stages in Figure 1 are desperately time-consuming (typically several hours) in large-scale real scenarios, hindering companies from building highly-responsive decision systems against fast-changing markets.

To achieve the ambition of rapid end-to-end optimization, both the performance bottleneck of instantiation and solving must be removed. In this paper, we propose two methods that fully address the efficiency of the two stages respectively, and encapsulate them as a new AMS, Grassland. For model instantiation, the motivation is to take advantage of both the sparsity of the data and the modern multiprocessor systems. While multiprocessor parallelism usually accompanies with an extra cost of communication and synchronization, we eliminate such a cost by a specially vectorized formulation that fully exploits the parallelism of data. In such a way we developed a model instantiation algorithm that is not only comparable with state-of-art AMSs in a single-threaded setting, but can also be accelerated in direct proportion to the number of processor cores. For model solving, we start from rolling horizon, a common business practice for decision making, to decompose a full model into a sequence of smaller models. While such a decomposition can lead to a significant loss of global optimality, we propose a heuristic of additional "aggregated master problem" to capture global optimality, so as to minimize the loss. We also integrate our proposed heuristics with an existing one to further improve the performance.

Our major contributions can be summarized as follows:

- Proposing a principled approach to stress the efficiency of model instantiation for optimization problems with linear constraints, by exploiting both the sparsity and parallelism of data, represented as a new algebraic modeling language with corresponding instantiation algorithms.
- Proposing Guided Rolling Horizon, a new decomposition heuristics that accelerates solving exponentially for mathematical optimization models with sequential structure with an acceptable loss of optimality, which can also work with other heuristics (Guided FRH) for better performance.
- Encapsulating our approaches as a new AMS, Grassland, and conducting extensive experiments on both classical LP/MIP models and real enterprise production planning scenarios.
This system has been deployed in Huawei's supply chain management scenario, especially for production planning simulation of our planning department. By reducing the end-to-end optimization pipeline from hours to 3-5 minutes, Grassland achieves our ambition of rapid optimization in real industrial scenarios, playing an essential role in the production plan decision making of Huawei against highly dynamic supply-demand environment.


## 2 RELATED WORK

Algebraic modeling systems/languages is a relatively mature field starting from the late 1970s, with many sophisticatedly designed
open-source and commercial software available. They can be divided into three categories. 1) Classical standalone systems with particular modeling language syntax, such as AMPL[10], GAMS[5] and ZIMPL[14]. They are usually more efficiently developed and easier to use for non-programmers; 2) Standalone modeling packages based on specific programming languages, such as YALMIP [15] for MATLAB, Pyomo [13] for Python, and JuMP[16] for Julia. Their modeling efficiency usually lies on the host language. 3) The modeling API as an integrated part of several solver engines, such as Gurobi[11] and Mosek[3], which usually provides solver-specific features.

For the model instantiation process, while most of the aforementioned AMSs stress the universality of modeling, hardly any of them take the efficiency issue (e.g., parallelization and vectorization) into full consideration, especially for large-scale scenarios with millions of variables. One may argue that such efficiency issue could be leaved to practitioners. However, it is not practical for hundredpage documented complex optimization models to be manually analyzed line-by-line for instantiation efficiency, as if manual calculation of gradients is not practical for complex deep learning models (although neither of them contains theoretical difficulties!). For complex models, principled design must be given to achieve practical implementation efficiency. While machine learning communities benefit enormously from such principledly designed frameworks such as TensorFlow[1] and PyTorch[17] , industrial optimization practitioners call for an analogous framework that boost the instantiation of highly complex business optimization models.

For decomposition methods of large-scale model solving, most of the current literature focuses on lossless decomposition such as Benders decomposition or column generation[4], in which the global optimality is guaranteed. However, such methods require a special block structure that is not easily satisfied in realistic complex scenarios. It is common that one or more types of constraints break the block structure, and blind use of decomposition methods will usually lead to performance degradation or even infeasibility. However, compared with feasibility, strict global optimality is not so crucial in most of the application scenarios. A fast-generated good solution is usually more appealing than a time-exhausting perfect one, which is the application foundation of heuristics methods. Existing lossy decomposition methods like forward rolling horizon [8] include a simple heuristics that aggregate future information to speedup the solving. In this way strict global optimality is slightly satisfied to exchange for more potential of acceleration. A detailed introduction is provided in Section 4.2.

## 3 EFFICIENT INSTANTIATION OF LINEAR CONSTRAINTS IN OPTIMIZATION

To design a rapid AMS, the first challenge is the efficiency of model instantiation for the emerging large-scale applied scenarios. In this section, we describe the model instantiation problem of mathematical optimization, and propose a general scheme to instantiate large-scale linear constraints efficiently, with both sparsity and parallelism taken into account.

### 3.1 Preliminaries

In this section, we give a brief introduction to mathematical optimization. A mathematical optimization (programming) problem can be represented in the following way: given a function $f: S \rightarrow \mathbb{R}$ from a set $S$ to real numbers, find an element $x_{0} \in S$ so that $\forall x \in S$, $f\left(x_{0}\right) \leq f(x)$ (minimization) or $f\left(x_{0}\right) \geq f(x)$ (maximization). Here $f, S$ and $x \in S$ are called objective function, feasible region and constraint respectively.

For the majority of optimization types in practice that can be solved efficiently, $S$ is a convex polytope. That is, to optimize an objective function subject to linear equality and inequality constraints, which can be expressed in canonical form as

$$
\begin{align*}
\min & f(\mathbf{x}) \\
\text { subject to } & \mathbf{A x}=\mathbf{b}, \mathbf{x} \geq \mathbf{0} \tag{1}
\end{align*}
$$

in which $A$ is called constraint matrix. $f(\mathbf{x})=\mathbf{c}^{T} \mathbf{x}$ for linear programming (LP) and $f(x)=\frac{1}{2} \mathbf{x}^{T} \mathbf{Q x}+\mathbf{c}^{T} \mathbf{x}$ for quadratic programming (QP). Mixed integer programming (MIP) adds additional integer constraints $x_{i} \in \mathbb{Z}, i \in Z$ on LP. Some special types of optimization contains nonlinear constraints such as conic and semidefinite optimization, which are out of this paper's scope.

Practically, while objective function $f(\mathbf{x})$ describe the business goal which is only a single expression, constraint matrix $\mathbf{A}$ can represent numerous business rules which contain millions of expressions that are much more time-consuming to instantiate. Therefore, we focus more on the efficiency of constraints instantiation in the following text.

### 3.2 Problem description and challenges

In this section, we describe the model instantiation problem of mathematical optimization, as well as its challenges in large-scale scenarios.
3.2.1 Problem description. While solver engines take canonical mathematical optimization problems as input, end users rarely write canonical problems directly. Instead, they develop symbolic representations of models in a human-readable language, which is called algebraic modeling language (AML). A symbolic representation of a model is a template of model without any concrete data. The place where concrete data should exist is represented by placeholders. When a symbolic model needs to be solved with given data, both the data and AML-based symbolic model are fed into AMS. AMS will compile the symbolic model, and fill the model with concrete data to generate a canonical representation for the solver engine. We name this process as model instantiation and conclude the input/output of this process as follows

Input (1) Symbolic representation of the model and (2) data.
Output Canonical representation of the model.
To show the procedure of model instantiation, we show an example of a simplified classical minimum-cost flow model. Consider a direct graph with a set $V$ of nodes and a set $E$ of edges, decision variable $x_{i, j}$ represents the amount of current flowing from node $i$ to node $j$. $s_{i}$ is the supply/demand at each node $i$. For each node $i$, flow out $\sum_{\{j \mid(i, j) \in E\}} x_{i, j}$ minus flow in $\sum_{\{j \mid(j, i) \in E\}} x_{j, i}$ must equal the supply/demand $s_{i}$. Every flow corresponds to a cost $c_{i, j}$, and the
model finds flows that minimize the total cost. The input/output for this model's instantiation is as follows

## Input:

- Symbolic representation of the model:

$$
\begin{align*}
& \min \sum_{(i, j) \in E} c_{i, j} x_{i, j} \\
& \text { subject to } \quad \operatorname{expr}_{i}=s_{i}, \quad \forall i \in V, \quad \mathbf{x} \geq \mathbf{0}  \tag{2}\\
& \text { in which } \quad \operatorname{expr}_{i}=\sum_{\{j \mid(i, j) \in E\}} x_{i, j}-\sum_{\{j \mid(j, i) \in E\}} x_{j, i} \tag{3}
\end{align*}
$$

in which $i, j$ are index placeholders denoting the index of expression expr and variable $x . V, E$ and $S$ are data placeholders whose value need to be specified in model instantiation process.

- Data: $V^{*}=\{1,2,3,4\}, E^{*}=\{(1,2),(2,4),(1,3),(3,4)\}, s^{*}=$ $[1,0,0,-1], c^{*}=\mathbf{1}$


## Output:

- Canonical representation of the model: (1) in which

$$
\left.\left.\begin{array}{rl}
\mathbf{c} & =[0,0,0,0,1,0,0,0,1,0,0,0,0,1,1,0
\end{array}\right]^{T}\right] \begin{array}{cccccccccccccccc}
0 & =\left[\begin{array}{ccccccccccccccc}
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & -1 \\
0
\end{array}\right] \\
\mathbf{b} & =[1,0,0,-1]^{T} \\
\mathbf{x} & =\left[x_{1,1}, \cdots, x_{4,1}, \cdots, x_{1,4}, \cdots, x_{4,4}\right]^{T}
\end{array}
$$

Feed $\mathbf{c}, \mathbf{A}$ and $\mathbf{b}$ into a LP solver engine and we can get the solution of decision variable $\mathbf{x}$. We will discuss the example output's generation process later in Section 3.3.2.
3.2.2 Challenges of model instantiation. While model instantiation seems trivial from a theoretical perspective (it can be easily achieved in polynomial time compared with model solving), the challenge is that, the time complexity can still be extremely high if we directly follow the literal meaning of mathematical expressions to instantiate models, especially for constraints. Consider the following general equality constraint

$$
\operatorname{expr}_{i_{1}, \cdots, i_{N}}=s_{i_{1}, \cdots, i_{n}}, \quad \forall i_{1} \in V_{1}, \cdots, i_{N} \in V_{N}
$$

The time complexity of direct instantiation is $O\left(\left|V_{1} \times \cdots \times V_{N}\right| \mid\right.$ expr $\left.\mid\right)$ in which $\mid$ expr $\mid$ is the computation cost to instantiate a single expression. For example, the time complexity of constraint (2)'s direct instantiation is $O(|V||E|)$ since by definition, we need to iterate every node $i(\forall i \in V$ in (2)), and search for edges whose heads or tails are equal to $i$ (sum operation $\sum_{\{j \mid(i, j) \in E\}}$ and $\sum_{\{j \mid(j, i) \in E\}}$ in (3)). This is clearly unbearable if we have a sparse graph with millions of vertices and edges.

In application scenarios, such a lack of principled, efficient model instantiation scheme results in serious performance issues. A surprising fact is that model instantiation costs similar or even more time than model solving in many large-scale, complex enterprise scenarios. While ad-hoc solutions may exist for specific kind of problems ${ }^{1}$, our aim is to develop an AMS which is efficient for

[^1]

Figure 2: The expression tree of expression (3), with global index placeholder $i$ and local index placeholder $j$.
general linear model instantiation, whose model definition can be in arbitrary forms.

### 3.3 A principled scheme of efficient linear constraint instantiation

In this section, we propose a principled scheme to instantiate linear constraints efficiently. The motivation is to design an AML with a corresponding instantiation algorithm that fully exploit the sparsity and parallelism of data.
3.3.1 A new AML with symbolic multidimensional expression. To represent a model symbolically with AML, the main task is to develop a symbolic representation of mathematical expressions in objective function and constraints. In our system, for a model with $K$ types of constraints (e.g., the balance constraint (3) is one type of constraint. A typical applied model can contain dozens of types of constraints), we formulate the symbolic representation as the following format:

$$
\begin{align*}
\min & \operatorname{expr}_{G_{o}}^{o} \\
\text { subject to } & \left(\operatorname{expr}_{G}^{i}, \operatorname{sign}_{i}, r h s_{G}^{i}\right), 1 \leq i \leq K \tag{4}
\end{align*}
$$

in which the $i$ th type of constraint is formulated as a triple ( $\operatorname{expr}_{G}$, sign, $\left.r h s_{G}\right)$. That is, a symbolic expression $\operatorname{expr} r_{G}$, a sign $(=, \geq$ or $\leq$ ), and a symbolic array $r h s_{G}$. For example, (2) can be represented as a triple $\left(\operatorname{expr}_{G},=, s_{G}\right)$ in which $G=(i)$. Here we assume that all constants are moved to the right hand side of the equation, and omit the superscript $i$ of $\operatorname{expr} r_{G}^{i}$ for simpler notation.

It is important to notice that the symbolic expression $\operatorname{expr}_{G}$ is multidimensional. That is, it represents a list of indexed expressions instead of a single one. For example, expr $r_{i}$ in (3) actually represents a list of expressions $\left[e x p r_{1}, \cdots\right.$, expr $\left.\left.\right|_{|V|}\right]$. Therefore, while a normal symbolic expression expr can be represented as an expression tree whose non-leaf and leaf nodes are operators and terms respectively, A multidimensional expression should be added with additional index placeholders $G$ to denote the index of expressions, represented as $\operatorname{expr}_{G}$. It consists of the following elements

- (Global) index placeholders $G=\left(g_{1}, \cdots, g_{N}\right)$, a tuple of one or more placeholders representing the index of expressions, defined at the root of the expression tree. Global index placeholders can appear anywhere in the expression tree. e.g., $G=(i)$ in (3), and the index placeholder $i$ also appears in the logical condition of two sum operators in (3). Here the bracket in $G=(i)$ indicates that $G$ is a tuple even if $N=1$.
- Operators (non-leaf nodes), receiving one or more nodes as operands. Some specific operators such as sum $(\Sigma)$ include local index placeholders that are only valid in the scope of these operators. E.g. the sum operator $\sum_{\{j \mid(i, j) \in E\}}$ in (3) receives $x_{i, j}$ as an operand, and include local index placeholder $j$ that is only valid in the scope of this sum operator. We will discuss the sum operator in detail later.
- Terms (leaf nodes), a variable with its corresponding coefficient. Their indices can be denoted by previously defined global and local index placeholders. e.g. $x_{i, j}$ in (3) is a term with variable $x_{i, j}$ and coefficient 1 .
and a example is shown in the left part of Figure 2.
In this work we mainly focus on the linear constrainted scenario that includes three operators: add $(+)$, subtract $(-)$ and sum $(\Sigma)$. While the add and subtract operators are trivial, we discuss the sum operator in detail. In our formulation, we number all sum operators in a expression, and sum operators take a fixed format: the $k$ th sum operator is assigned to a symbolic logical condition $\left(G_{k} \| L_{k}\right) \in S_{k} .{ }^{2}$ "\|" denotes the concatenation of tuples. For example, the 1 st $(k=1)$ sum operator $\sum_{(i, j) \in E}$ in (3) contains a logical condition $(i, j) \in E$. It consists of
- Global index placeholders $G_{k}=\left(g_{k_{1}}, g_{k_{2}}, \cdots\right)$, a tuple whose elements are the index placeholders from G. e.g., $G_{1}=(i)$ in $\sum_{(i, j) \in E}$.
- Local index placeholders $L_{k}=\left(l_{1}, l_{2}, \cdots\right)$, which is only valid in the scope of this sum operator. e.g., $L_{1}=(j)$ in $\sum_{(i, j) \in E}$.
- Data placeholder $S_{k}$ : A symbolic placeholder representing a set of fixed-sized index tuples. The size of each tuple is $\left|G_{k}\right|+$ $\left|L_{k}\right|$. e.g., edge set placeholder $E$ in $\sum_{(i, j) \in E}$, representing a set of direct edges (i.e., pair of node indices).
An example is shown in the right part of Figure 2.
3.3.2 An efficient model instantiation algorithm. To instantiate a multidimensional expression $\operatorname{expr}_{G}$ given data $S^{*}$, a simple way is to enumerate all possible combinations of global index placeholders $G$ (denoted as space $(G)$ ), and traverse through the expression tree for each combination to generate every single algebraic expression. This is usually the literal meaning of mathematical expressions. The detailed process is illustrated in Algorithm 1. For example, to instantiate (3) with given data $V^{*}$ and $E^{*}$, this expression's index placeholders is $G=(i)$, and all possible values of $i$ are space $(G)=$ $\{1,2,3,4\}$. Then we iterate all elements of space $(G)$. e.g., when $i=1$, we iterate $E$ to generate sub-expression $\sum_{(1, j) \in E} x_{1, j}=x_{1,2}+x_{1,3}$ and $\sum_{(j, 1) \in E} x_{j, 1}=0$, then the expression will be expr $r_{1}=x_{1,2}+$ $x_{1,3}$. In a similar way we get expr $r_{2}=x_{2,4}-x_{1,2}, \operatorname{expr}_{3}=x_{3,4}-$ $x_{1,3}$, expr $_{4}=-x_{2,4}-x_{3,4}$.

However, this approach is extremely exhaustive with time complexity $O\left(\left|\operatorname{space}(G) \| S^{*}\right|\right)$. In this section, we propose an efficient model instantiation algorithm based on the AML proposed in previous section, whose result is identical to Algorithm 1 but with $\left|O\left(S^{*}\right)\right|$ time complexity.

Lemma 3.1. For expression tree of $\operatorname{expr}_{G}$, without loss of generality, we assume that every path from a leaf node to the root will go through at least one sum operator.

[^2]```
Algorithm 1 Exhaustive model instantiation algorithm
    Input: Symbolic multidimensional expression \(\operatorname{expr}_{G}\), set data \(S^{*}\)
    Output: Constraint matrix \(A\)
    1: Initialize \(A\) as an empty matrix, with number of columns equal to number of
        variables
        for \(G^{*} \in \operatorname{space}(G)\) do
            Generate \(\operatorname{expr}_{G}^{*}\) by replacing symbolic index placeholders \(G\) and data place-
        holder \(S\) with concrete value \(G^{*}\) and \(S^{*}\) respectively in all nodes of \(\operatorname{expr} r_{G}\)
        Do an in-order traversal to expr \({ }_{G}^{*}\) to generate the algebraic expression. When
        visiting the \(i\) th sum operator, traverse through corresponding set data \(S_{i}^{*}\) in its
        logical condition.
            Add one row to \(A\) with the generated expression.
    end for
    return \(A\)
```

Proof. For the leaf node whose path to the root does not include any sum operator, we can insert a "dummy" sum operator $\sum_{G \in S_{i}}$ before the node with set data $S_{i}^{*}=\operatorname{space}(G)$. This operator does not contain any local index placeholders so will not change the result of Algorithm 1.

Lemma 3.2. Let $I($ node $)=\{i \mid i t h$ sum operator is on the path from node to root\}. Without loss of generality, we assume that for every leaf node of the expression tree, $\left\{g \mid g \in G_{j}, j \in I(\right.$ node $\left.)\right\}=G$.

Proof. From Lemma 3.1 we know that $I($ node $) \neq \emptyset$. If $\exists g^{\prime} \in G$ so that $g^{\prime} \notin\left\{g \mid g \in G_{j}, j \in I(\right.$ node $\left.)\right\}$, we select one sum operator $\sum_{\left(G_{i} \| L_{i}\right) \in S_{i}}$ on $I($ node $)$ and expand $G_{i}$ to $G_{i} \| g^{\prime}$. For set data $S_{i}^{*}$, we replace each tuple data $\left(g_{i_{1}}^{*}, \cdots, g_{i_{N}}^{*}, l_{i_{1}}^{*}, \cdots\right)$ to a set of expanded tuple $\left\{\left(g_{i_{1}}^{*}, \cdots, g_{i_{N}}^{*}, g_{j}^{*}, l_{i_{1}}^{*}, \cdots\right) \mid g_{j}^{*} \in \operatorname{space}\left(g_{j}\right)\right\}$. In this way we enumerate all possible value of $g_{j}$ for every tuple data in $S_{i}^{*}$, so will not change the result of Algorithm 1.

With Lemma 3.1 and Lemma 3.2, we propose a model instantiation algorithm. Different from Algorithm 1 that fixes the value of all global index placeholders and traverses the expression tree for |space $(G) \mid$ times, this algorithm traverses the expression tree only once, and records the corresponding value of index placeholders dynamically as "context information" when traverse through concrete set data $S_{i}^{*}$ of $i$ th sum operator. When the leaf node (term) is reached, all the index placeholders in the term is replaced by the actual value recorded in the context information. Meanwhile, the actual value of all global index placeholders $G$ in the context is snapshotted and attached to the index-replaced term. When the traverse process is finished, we aggregate terms with the same global index. The detailed algorithm is shown in Algorithm 2.

For example, to instantiate (3), we traverse the expression tree of (3) in Figure 2. When we arrive at the first sum operation $\sum_{(i, j) \in E}$, we iterate the value $(i, j) \in E$ and record the context information (e.g., record $i=1, j=2$ for the first edge). When we reach the leaf node $x_{i, j}$, we replace the index placeholder with the corresponding value recorded in context information (e.g., we get $x_{1,2}$ ), and attach the actual value of all global index placeholders to the term (e.g., attach $i=1$ to $x_{1,2}$, represented as $\left(1, x_{1,2}\right)$ ). When we finish the traverse process, we will get $\left(1, x_{1,2}\right),\left(2, x_{2,4}\right),\left(1, x_{1,3}\right),\left(3, x_{3,4}\right)$, $\left(2,-x_{1,2}\right),\left(4,-x_{2,4}\right),\left(3,-x_{1,3}\right),\left(4,-x_{3,4}\right)$. By aggregating terms with the same global index, we will get the same result as Algorithm 1.

Proposition 3.1. The outputs of Algorithm 1 and Algorithm 2 are identical given the same input.

```
Algorithm 2 Efficient model instantiation algorithm
Input: Symbolic multidimensional expression \(\operatorname{expr}_{G}\), set data \(S^{*}\)
Output: Constraint matrix \(A\)
    1: procedure Replace(node, context) // node is a leaf node (term)
    2: Replace the index placeholders of coefficient and variables with concrete values
    in context, and return the replaced term
    end procedure
    procedure Replace( \(G\), context) // G is a tuple of placeholders
        Replace the global index placeholders in \(G\) with concrete values in context
    and return
    end procedure
    procedure Iterate(node, context)
            if node is a leaf node then
                terms \(\leftarrow(\operatorname{Replace}(\) node, context), Replace \((G\), context))
            else if node is the \(i\) th sum operator then
                    terms \(\leftarrow\) emply list
                    Retrieve the concrete data \(S_{i}^{*}\) of the \(i\) th sum operator from \(S^{*}\)
                    Filter all \(\left(g_{1}^{*}, \cdots, g_{N}^{*}\right) \in S_{i}^{*}\) with the condition that \(g_{j}^{*}=\operatorname{context}\left(g_{j}\right)\) if
    \(g_{j}\) appears in the mapping key of the context
                for \(\left(G_{i}^{*}, L_{i}^{*}\right) \in S_{i}^{*}\) do
                    context' \(\leftarrow\) AddMAPping (context, \(\left.G_{i} \rightarrow G_{i}^{*}, L_{i} \rightarrow L_{i}^{*}\right)\)
                    terms \(\leftarrow\) terms \(\|\) ITERATE(child, context')
                end for
            else if node is an add operator then
                terms \(\leftarrow\) Iterate(left, context) || Iterate(right, context)
            else if node is a sub operator then
                terms \(\leftarrow \operatorname{Iterate}(\) left, context) \(\|\)-Iterate(right, context)
                end if
                return terms
    end procedure
    Initialize \(A\) as an empty matrix sized \(\mid\) constraints \(|\times|\) variables \(\mid\).
    Initialize context as an empty mapping from symbolic index placeholder to con-
    crete value.
    terms \(\leftarrow\) Iterate(root node of expr \(r_{G}\), context)
    for (term, \(G^{*}\) ) \(\in\) terms do
        Map \(G^{*}\) and variable to matrix index row and col
        \(A[\) row, col \(]=\) coefficient of variable
    end for
    return \(A\)
```

Proof. To simplify the demonstration we omit the coefficient in all terms, which can be treated similarly to the variables.
$\Rightarrow$ : In Algorithm 1, assume there is a variable $x_{G_{x}^{*}, L_{x}^{*}}$ in expression $\operatorname{expr}_{G^{*}}$, from Lemma 3.2 we know that the union of all $G_{i}$ in $I\left(x_{G_{x}^{*}, L_{x}^{*}}\right)$ equals to $G$. Without loss of generality we let $I\left(x_{G_{x}^{*}, L_{x}^{*}}\right)=1, \cdots, M$, then follow the depth-first iteration of Algorithm 1 , we can find a sequence $\left(G_{1}^{*}\left\|L_{1}^{*} \in S_{1}^{*}, \cdots, G_{M}^{*}\right\| L_{M}^{*} \in S_{M}^{*}\right)$ so that $\bigcup_{i=1}^{M} G_{i}=G$ and $\bigcup_{i=1}^{M} G_{i}^{*}=G^{*}$. For Algorithm 2, we can follow the same sequence in the depth-first iteration and accumulate the mapping $G_{i} \rightarrow G_{i}^{*}, L_{i} \rightarrow L_{i}^{*}$ in context. Therefore when the leaf node is finally reached, the context will contain $G \rightarrow G^{*}$, then by the Replace procedure we get variable $x_{G_{x}^{*}, L_{x}^{*}}$ in expression $\operatorname{expr}_{G^{*}}$ in Algorithm 2.
$\Leftarrow$ : In Algorithm 2 assume there is a variable $x_{G_{x}^{*}, L_{x}^{*}}$ in expression $\operatorname{expr}_{G^{*}}$, the context information stores mapping $G \rightarrow G^{*}$ when Iterate reached the leaf node, then following Lemma 3.2 we also have a sequence ( $G_{1}^{*}\left\|L_{1}^{*} \in S_{1}^{*}, \cdots, G_{M}^{*}\right\| L_{M}^{*} \in S_{M}^{*}$ ) so that $\bigcup_{i=1}^{M} G_{i}=G$ and $\bigcup_{i=1}^{M} G_{i}^{*}=G^{*}$. Thus when $\operatorname{expr}_{G}$ is fixed into $\operatorname{expr}_{G^{*}}$ in Algorithm 1, we can follow the same sequence and get variable $x_{G_{x}^{*}, L_{x}^{*}}$ in expression $\operatorname{expr}_{G^{*}}$.
3.3.3 Parallelization of the model instantiation algorithm. While Algorithm 1 is extremely exhaustive, it is easy to be paralleled by simply letting each worker instantiate a partition of concrete index set space $(G)$. In this section we show that Algorithm 2 can also be fully paralleled by transiting index partition to data partition.


Figure 3: An example of data partition.
Given a set data $S^{*}$, its corresponding index placeholders $G$ and a partition of all possible values of $k$ th index placeholder $g_{k} \in G$ (denoted as $\left\{P_{1}, \cdots, P_{M}\right\}, \bigcup_{i} P_{i}=\operatorname{space}\left(g_{i}\right)$ ), a data partition of set data $S^{*}$ over $k$ th index $g_{k}$ is to partition $S^{*}$ to $M$ subset $S_{1}^{*}, \cdots, S_{M}^{*}$, so that $S_{i}^{*}=\left\{\left(g_{1}^{*}, \cdots, g_{N}^{*}\right) \mid\left(g_{1}^{*}, \cdots, g_{N}^{*}\right) \in S^{*}, g_{k}^{*} \in P_{i}\right\}$. That is, the $i$ th subset of the data only contains the index tuple whose $k$ th element is in $i$ th partitioned index subset $P_{i}$. An example is shown in Figure 3.
Then we have the following proposition:
Proposition 3.2. Running Algorithm 1 on a subset of all possible global index value $\left(g_{1}^{*}, \cdots, g_{N}^{*}\right) \in \operatorname{space}(G)$ with condition $g_{k}^{*} \in P$, is equivalent to running Algorithm 2 on a subset of data $\left(g_{1}^{*}, \cdots, g_{N}^{*}\right) \in$ $S^{*}$ with condition $g_{k}^{*} \in P$.

Proof. The proof is similar to Proposition 3.1 with some details on sequence sharing between algorithms. For $\Rightarrow$, the sequence $\left(G_{1}^{*}\left\|L_{1}^{*} \in S_{1}^{*}, \cdots, G_{M}^{*}\right\| L_{M}^{*} \in S_{M}^{*}\right)$ can still be transferred to Algorithm 2 whose data $S^{*}$ is filtered by condition $g_{k}^{*} \in P$, since $\operatorname{space}(G)$ is also filtered in Algorithm 1 to make sure only $g_{k}^{*} \in P$ will appear in the sequence. For similar reason $\Leftarrow$ holds.

With Proposition 3.2 we can fully parallelize Algorithm 2. For example, to distribute the instantiation process of (3) equally to two workers, by applying Proposition 3.2, we can do data partition on $E^{*}$ as $E_{1}^{*}=\{(1,2),(2,4),(1,3)\}, E_{2}^{*}=\{(1,2)\}$ for worker 1, $E_{1}^{*}=$ $\{(3,4)\}, E_{2}^{*}=\{(2,4),(1,3),(3,4)\}$ for worker 2. Applying Algorithm 2 , worker 1 will generate $\left(1, x_{1,2}\right),\left(2, x_{2,4}\right),\left(1, x_{1,3}\right),\left(2,-x_{1,2}\right)$, worker 2 will generate $\left(3, x_{3,4}\right),\left(4,-x_{2,4}\right),\left(3,-x_{1,3}\right),\left(4,-x_{3,4}\right)$. It is easy to check that the results are identical.

## 4 SEQUENTIAL DECOMPOSITION OF LARGE-SCALE OPTIMIZATION

The second challenge for a rapid AMS is the solving time. Extremely large-scale mathematical optimization models usually take huge amount of time to solve, and the scale of data may even bump up due to sudden business need (e.g., big promotion), resulting in potential risk of timeout or even unsolvability. In this section, we introduce a lossy decomposition method (Guided FRH) for massive sequential models. While the feasibility of solutions is strictly maintained, the decomposition methods make a trade-off between optimality and efficiency.

### 4.1 Preliminaries: Sequential Decision Making and Rolling Horizon

Sequential (or dynamic) decision making widely exists in applied scenarios. For example, we may need to decide the production amount of specific items in a range of dates, in which the prior
decisions will influence successive ones. More formally, for a sequential model of linear constraints with $T$ periods, its decision variables $\mathbf{x}$ can be divided into $T$ row vectors $\mathbf{x}_{1}, \cdots, \mathbf{x}_{T}$, so that the constraints can be formulated as

$$
\begin{array}{ll}
\mathbf{A}_{1} \mathbf{x}_{1}^{T} & =\mathbf{b}_{1} \\
\mathbf{A}_{2}\left[\mathbf{x}_{1}, \mathbf{x}_{2}\right]^{T} & =\mathbf{b}_{2} \\
\ldots & \\
\mathbf{A}_{T}\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{T}\right]^{T} & =\mathbf{b}_{T}
\end{array}
$$

which indicates that the constraint matrix has a block triangular structure. Here we assume the decision variables share the same semantic meaning in each period, and use $x_{t}^{i}$ to denote the $i$ th variable in period $t$. We refer to [4] for a detailed introduction.

To make sequential decisions in a dynamic environment, a common business practice is rolling horizon (RH)[19]. That is, we do planning in a relatively long time window (planning horizon) using the latest available information, and only accept the generated decisions in the first several time steps (re-planning periodicity). When the re-planning periodicity passed, we start a new planning horizon with updated environment information, and repeat the above procedure.

From this perspective, compared with the decision in the replanning periodicity that will be actually applied, the planning after the re-planning periodicity is more likely a "simulation-based guidance" to reach the global optimum. That is, although the decisions after the re-planning periodicity is never actually executed, we assume that they will be executed on a simulation basis, so that the global optimum in a longer range is considered.

In our large-scale scenario, the size of the planning horizon is limited due to the scalability of solver engines. To enlarge the range of future information involved in the optimization and reach global optimum in a longer time window, we adopt the idea of "guidance" in rolling horizon, but with a more computational efficient approach.

### 4.2 Forward Rolling Horizon

To begin with, we introduce a simple method, Forward Rolling Horizon (FRH) [8] to illustrate the basic idea. For a large sequential model $P$ with length $T$, we divide the model into $h$ sub-models $P_{1}, \cdots, P_{h}$. Each sub-model $P_{i}$ starts at sequence period $t_{i}$ and ends at period $t_{i+1}-1$. The sub-problems are solved in a sequential manner so that the model $P_{i}$ can take advantage of the solutions of $P_{1}, \cdots, P_{i-1}$. To guide each sub-model $P_{i}$ towards the global optimum, we aggregate the future information from period $t_{i+1}$ to the last period into $M$ periods ( $M=1$ by default), and attach it to the end of the sub-model. Therefore, except for the last sub-model, each sub-model consists of $t_{i+1}-t_{i}+M$ periods. The FRH optimization procedure is shown in Figure 4a.

### 4.3 Guided Rolling Horizon

In this section, we propose Guided Rolling Horizon (Guided RH), which decompose a raw sequential problem into a main problem and a sequence of sub-problems, with all constraints strictly satisfied and global optimality largely preserved.


Figure 4: Forward rolling horizon and proposed decomposition methods for large-scale mathematical optimization.

For a raw sequential problem with $T$ periods and $N$ decision variables in each period, the Guided RH optimization plan consists of three steps:

Data aggregation Aggregate all feed-in data from $T$ periods to $h$ periods, in which $h$ is the number of sub-models. Every data element at period $i$ in the aggregated data represents the data elements from period $t_{i}$ to $t_{i+1}-1$ in the original data. E.g., aggregate data from daily to weekly.
Master problem solving Construct a "shrunk" master problem with $h$ periods using the aggregated data, then solve it. The master problem provides a high-level, soft guidance to the actual rolling-horizon procedure afterwards, therefore all the integer constraints can be omitted. Similar to data aggregation step, every solution element $z_{k}^{i}$ at period $k$ in this master problem represents the solution elements $x_{t_{k}}^{i}, \cdots, x_{t_{k+1}-1}^{i}$ in the original problem.
Sub problems solving To transit the master problem's aggregated solutions to an applicable solutions, we apply the rolling horizon method on the original problem, with additional guided constraints and objectives to make the two solutions aligned. In each horizon $k$ (from period $t_{k}$ to period $t_{k+1}-1$ ), we add soft constraints to each variable group along period axis, to minimize the gap between applicable solutions $x_{t_{k}}^{i}, \cdots, x_{t_{k+1}-1}^{i}$ and the aggregated solution $z_{k}^{i}$. i.e.,

$$
\begin{array}{ll}
\min & o b j+\sum_{i=1}^{N} \lambda_{k}\left(u_{k}^{i}+v_{k}^{i}\right) \\
\text { s.t. } & \mathbf{A}_{k}\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{k}\right]^{T}=\mathbf{b}_{k} \\
& \sum_{j=t_{k}}^{t_{k+1}-1} x_{j}^{i}-z_{k}^{i}=u_{k}^{i}-v_{k}^{i}, \quad \forall i \in\{1, \cdots, N\} \\
& u_{k}^{i} \geq 0, v_{k}^{i} \geq 0, \quad \forall i \in\{1, \cdots, N\}
\end{array}
$$



Figure 5: The fine-tuning procedure. The grey variables are fixed while the green variables are to be re-optimized. State variables whose value are determined by other variables keep free in the whole sequence.
in which $\mathbf{x}_{1}, \cdots, \mathbf{x}_{k-1}$ are already solved and fixed in previous sub-problem solving. obj is the original objective of the problem, $u_{i}$ and $v_{i}$ are auxiliary variables that help minimize the L1 loss $\sum_{i}\left|\sum_{j=t_{k}}^{t_{k+1}-1} x_{j}^{i}-z_{k}^{i}\right| \cdot \lambda_{k}$ is the weight controlling to what extent should the applicable solution be aligned with the aggregated solution in horizon $k$. It is optional to re-solve the master problem with current horizon's solution fixed to reduce the cumulative error between master and sub-problems. It is also possible for sub-problems to overlap with each other. For many real scenarios, only the solution of first $K$ periods is needed, in this case we can stop after solving $l$ horizons so that $t_{l+1}>K$.

The procedure of Guided RH is shown in Figure 4 b .

### 4.4 Guided FRH

While both FRH and Guided RH guide rolling horizon towards long-term global optimality, there is no conflict between them. Therefore, we can further improve the optimality by replacing the rolling horizon solving procedure of Guided RH with FRH so as to combine the two methods. The procedure of Guided FRH is shown in Figure 4c.

### 4.5 Fine-tuning of approximated solutions

When we get an approximated solution $\mathbf{x}_{1}, \cdots, \mathbf{x}_{T}$ of a sequential model with the above decomposition methods, it will be beneficial if we can do some fine-tuning to the solution to further boost the global optimality. A possible method is to re-optimize the variable in the first $k$ th periods, with other variables fixed. Since the solution is already a feasible one, the re-optimized solution will not be worse than the original solution. Notice that there might be some "state variables" whose value are fully determined by other variables. These variables will keep free in the whole sequence, so that the change of previous decisions can still propagate to the later periods. The procedure of fine-tuning is shown in Figure 5.

## 5 EVALUATION

In this section, we first conduct offline benchmarks against current state-of-the-art AMSs on model instantiation efficiency, and then deployed Grassland in Huawei's production planning scenario for nearly half a year. By providing near-real-time production planning simulation, Grassland plays an essential role in Huawei's supply chain management against highly dynamic supply-demand environment.

|  | P-Median | Offshore <br> Wind Farming | Food <br> Manufacture I |
| :---: | ---: | ---: | ---: |
| Gurobi Py API | 410.20 | 533.71 | 744.39 |
| JuMP | 278.08 | 169.08 | 789.86 |
| ZIMPL | 174.00 | 400.47 | 399.16 |
| AMPL | 15.94 | 17.71 | 31.65 |
| Grassland (S) | 35.91 | 18.85 | 80.83 |
| Grassland (M) | $\mathbf{2 . 0 9}$ | $\mathbf{1 . 6 7}$ | $\mathbf{5 . 2 8}$ |

Table 1: Model Instantiation Benchmark. Total time (in seconds) to process the model definition and produce the output file in CPLEX LP format.

### 5.1 Offline Model Instantiation Benchmark

Before online experiment, we first benchmark our system against several widely-used modeling software on some typical mathematical optimization problems, to evaluate the scalability and effectiveness of proposed model instantiation method.
5.1.1 Test Problems. We select three typical problems from JuMP [16] and Gurobi modeling examples [12].

- P-Median: This model is used in $[13,16]$ to compare the modeling efficiency of Pyomo and JuMP with other AMSs.
- Offshore Wind Farming: This model in Gurobi modeling example minimize the cost of laying underwater cables to collect electricity produced by an offshore wind farm.
- Food Manufacture I: This model in Gurobi modeling example is a blending problem. Multiple raw materials are combined in a way to achieve the lowest cost.
5.1.2 Compared Modeling Softwares. In the following experiments, we compare the following modeling software. Modeling software implemented in interpreted languages such as Pyomo [13] are not included due to lack of efficiency in the prior benchmark works [13, 16].
- Gurobi Modeling API [11]: The Python interface of Gurobi, which provides the standard implementation of Gurobi modeling examples.
- ZIMPL [14]: The modeling language in SCIP Optimization Suite, written in plain C.
- JuMP [16]: A modern AML implemented in Julia, which is reported to achieve high efficiency comparable to commercial products such as AMPL, with the benefits of remaining in a high-level modern language.
- AMPL [10]: A leading, widely-used commercial AMS.
- Grassland (Single-threaded): The implementation of our proposed method in Section 3.3 with only one thread.
- Grassland (Multi-threaded): The implementation of our proposed method, including parallelization described in Section 3.3.3 ( 64 threads).
5.1.3 Results. The benchmark result on all the problems is shown in Table 1. It is shown that Grassland achieve absolute predominance over all other modeling software. While single-threaded Grassland already achieve $4-5 \mathrm{x}$ speedup over the fastest opensource modeling software (ZIMPL) and comparable with the leading commercial software (AMPL), multi-threaded Grassland further achieves $6-10 \mathrm{x}$ speedup over AMPL.


Figure 6: Offline model instantiation benchmark on (a) P-Median (b) Offshore Wind Farming (c) Food Manufacture I.


Figure 7: Online experiment results. (a) Comparison between baseline and different decomposition methods on time cost and optimality. (b) Same as (a) with fine-tuning of first 20 periods. (c) Comparison of model instantiation efficiency.

We also tested the software on different scales of models. The result is shown in Figure 6. It is shown that Grassland has superior performance over all scale of models.

### 5.2 Online Experiment

5.2.1 Background: Production planning and supply-demand analysis. Production planning is the planning of production activities to transform raw materials (supply) into finished products, meeting customer's order (demand) in the most efficient or economical way possible. It lies in the core of manufacturing companies' supply chain management, directly influencing the profit and customer satisfaction. Mathematical optimization is a mainstream method for production planning, and [18] provides a comprehensive introduction to this. As a world-leading electronic manufacturer, Huawei provides more than ten thousand kinds of end products, with even much more kinds of raw materials and intermediate assemblies, to satisfy millions of demands from all over the world. Modeling in such a large-scale scenario involves millions of variables and constraints.

While mathematical optimization can deliver near-optimal production planning solution for a certain input of supply and demand, The demand and supply itself is always changing with high uncertainty due to external dynamic factors. Our planning department needs to react quickly to such changes to ensure business continuity, which is called supply-demand analysis. To support the analysis, a crucial process is production planning simulation which returns the final planning result (e.g., fulfillment rate) for a certain input of supply and demand, helping planners evaluate and improve their
analysis. For example, when several raw materials are suddenly unavailable, planners may try different ways to increase the supply of alternative materials and run production planning simulation for each of them, adjust the supply iteratively to increase the fulfillment rate of related end products, and get the final supply adjustment decision.
5.2.2 Dataset and Compared Methods. The dataset is from real production environment which consists of all needed manufacturing data in 78 weeks (one week per period). The instantiated model consists of 5,957,634 variables, 3,443,465 constraints and 30,366,971 nonzero elements in the constraint matrix. In the following experiments, we compare the result before and after the application of Grassland in the production planning simulation scenario. The baseline is the original method before the deployment of Grassland, which is quite standard: ZIMPL is used for modeling and instantiation, and the instantiated full model is directly solved without any decomposition. The decomposition methods in section 4 (RH, FRH, Guided RH, Guided FRH) are tested separately. Sequence length $T=78$, number of submodels $h=8$, fine-tuning periods $k=20$. All models are solved via Mosek $\mathrm{ApS}[3] .{ }^{3}$
5.2.3 Results. The main result is shown in Figure 7a. Our proposed methods can achieve significant acceleration (15-35x faster) than baseline method, while the feasibility is strictly maintained and

[^3]the loss of objective is small. For Guided FRH, it can achieve 15 x acceleration with only $0.36 \%$ of the objective loss. Practically, such a tiny optimality loss does not cause sensible issues $(0.1 \% \sim 0.2 \%$ fluctuation of fulfillment ratio), especially considering that multiple source of more dominant error exist in complex business models such as prediction and approximation error.

The fine-tuning result is shown in Figure 7b in which the first 20 periods of the problem is re-optimized following Section 4.5. It is shown that fine-tuning can significantly narrow the gap between the objective value of decomposition methods and the optimal one.

Additionally, for instantiation efficiency in online scenarios, we compare Grassland with two legacy systems that we previously developed and deployed in online environment, based on ZIMPL and Gurobi Modeling API respectively. We use the number of periods involved to control the model size. The result is shown in Figure 7c. It is shown that the result is aligned with offline benchmarks in Figure 6.

## 6 CONCLUSION

In this paper, we propose Grassland, an algebraic modeling system that is efficient in large-scale mathematical optimization scenarios, including a parallelized instantiation scheme for general linear constraints, and a lossy sequential decomposition method that accelerates large-scale model solving exponentially. We perform both offline benchmarks and online deployment in Huawei's production planning scenario. The results demonstrate the significant superiority of Grassland over strong baselines.

## APPENDIX

## A IMPLEMENTATION DETAILS

As an algebraic modeling system, Grassland is implemented in five layers:

Modeling API (AML) The grassland modeling API is implemented as a Python package grassland (gl).
Intermediate representation (IR) layer This layer plays as a bridge between Modeling API and highly-efficient c++ backend. Models defined by Grassland modeling API is translated into a unified, JSON-based intermediate representation with six components (variables, constants, index placeholders, expression graphs, constraints and bounds).
Decomposition layer Implements four decomposition methods in section 4 (RH, FRH, Guided RH, Guided FRH).
Model instantiation layer Implements the parallelized model instantiation scheme in Section 3.3.2 and Section 3.3.3. The parallelization is implemented by multi-threaded programming. Due to the extreme efficiency of our proposed method, even the float-to-string conversion becomes a significant bottleneck. Here we apply Ryu [2] to accelerate the conversion.
Solver layer Calls different solver engine to solve the instantiated model or sub-model and return back the solution.

## B INTEGRATION WITH ROUNDING PROCEDURE

In applied optimization pipeline, there usually exist some integer constraints for decision variables. To achieve this, a mixed integer

| Test Problem | \#(variables) | \#(constraints) | \#(nonzeros) |
| :---: | ---: | ---: | ---: |
| P-Median | $5,050,000$ | $5,000,164$ | $15,050,000$ |
| Offshore Wind Farming | $4,170,120$ | $4,220,120$ | $10,425,300$ |
| Food Manufacture I | $5,006,394$ | $14,974,502$ | $39,932,750$ |

Table 2: The basic statistics of the benchmark test problems in Table 1.
programming model or an external heuristics rounding procedure may apply. However, the efficiency of rounding procedure can heavily rely on the scale of the model, as well as the number of integer constraints. With the above decomposition methods, we can round the variables at the same time when we solve each sub-model. Since the sub-model is significantly smaller than the original one, the rounding procedure will also be largely accelerated.

## C TEST PROBLEMS FOR MODEL INSTANTIATION

The basic statistics of the benchmark test problems in Table 1 is listed in Table 2. Problem data is randomly generated while maintaining the feasibility, and we control the size of the data to generate different scale of models in Figure 6.

## D EXPREIMENTAL SETTING

For offline model instantiation benchmark, all benchmarks are run on a server with 32 -core (64-thread) CPUs and 192GB memory. The output format of all constructed problems is set to CPLEX LP format (a standard LP/MIP format that is supported by most of the mathematical solvers). The identity of all constructed problems by different softwares on small and medium size are checked by actual solving with Gurobi with the same optimized objective value, and checked by static comparison script for extremely large size that cannot be directly solved in reasonable time. MIP problems are relaxed into LP in the identity checking process to extend scalability. For multi-threaded Grassland, the size of thread pool is set to 64 . The basic information of the benchmark problems are shown in Table 2.

For online experiment. all experiments are run on a Huawei cloud server with 32-core (64-thread) CPUs and 256GB memory. All models are solved via Mosek ApS [3]. For production planning problem, the sequence length $T=78$ and we use number of submodels $h=8$, fine-tuning periods $k=20$ for all decomposition methods.

Software version:

- Gurobi Modeling API (7.5.1, released in Jul 2017) ${ }^{4}$
- ZIMPL (3.4.0, released in June 2020)
- JuMP (0.21.3, released in June 2020)
- AMPL (20200810, released in Aug 2020)


## E EXPERIMENTAL RESULT

The detailed result of online experiment is shown in Table 3 and Table 4.

[^4]| Method | Model <br> instantiation timeSolving <br> time | Total <br> time | Objective <br> $\left(\times 10^{9}\right)$ |  |
| :---: | ---: | ---: | ---: | ---: |
| Baseline | 857.00 | 3135.16 | 3992.16 | 5.15626 |
| RH | 4.94 | 40.80 | 45.74 | $(+7.45 \%) 5.54017$ |
| FARH | $8.54^{*}+6.78$ | 67.68 | 83.00 | $(+0.81 \%) 5.19818$ |
| G-RH | $28.84^{* *}+8.79^{*}+6.20$ | 73.31 | 117.14 | $(+2.11 \%) 5.26513$ |
| G-FARH | $30.99^{* *}+11.07^{*}+7.43$ | 167.62 | 217.11 | $\mathbf{( + 0 . 3 6 \% )} 5.17477$ |

Table 3: Sequential decomposition benchmark on demandsupply analysis problem. In "Model instantiation time" column, time marked with "*" is the time for data compression, time marked with "**" is the time to generate guided constraints and objectives.

| Method | Model <br> instantiation time | Solving <br> time | Total <br> time | Objective <br> $\left(\times 10^{9}\right)$ |
| :---: | ---: | ---: | ---: | ---: |
| Baseline | 857.00 | 3135.16 | 3992.16 | 5.15626 |
| RH | 8.16 | $150.35^{\dagger}+40.80$ | 199.31 | $(+3.52 \%) 5.33775$ |
| FARH | $8.54^{*}+9.19$ | $147.11^{\dagger}+67.68$ | 232.52 | $(+0.78 \%) 5.19651$ |
| G-RH | $28.84^{* *}+8.79^{*}+9.29$ | $138.11^{\dagger}+73.31$ | 258.34 | $(+1.26 \%) 5.22104$ |
| G-FARH | $30.99^{* *}+11.07^{*}+9.93$ | $116.41^{\dagger}+167.62$ | 336.02 | $(+\mathbf{0 . 3 3 \% )} 5.17321$ |

Table 4: Fine-tuning benchmark on demand-supply analysis problem. In "Solving time" column, time marked with " $\dagger$ " is the extra solving time for fine tuning.

## F PRODUCTION PLANNING MODEL

While real-world production planning models are complex with lots of variants for different scenarios, here we show a self-contained, simplified version with only three types of constraints. We refer to [18] for a detailed introduction.

$$
\begin{align*}
& \min \quad \sum_{t, p, i} C_{t, p, i}^{m} m_{t, p, i}+\sum_{t, p, i} C_{t, p, i}^{x} x_{t, p, i}+\sum_{t, p, i} C_{t, p, i}^{p u r} p u r_{t, p, i} \\
& +\sum_{t, p, i, i^{\prime}, j} C_{t, p, p, i, i^{\prime}, j}^{r p} r_{t, p, i, i i^{\prime}, j}+\sum_{t, p, i, j} C_{t, p, i, j}^{r} r_{t, p, i, j} \\
& \text { s.t. } \quad \text { inv }_{t, p, i}=\text { inv }_{t-1, p, i}+\text { inbound }_{t, p, i}-\text { outbound }_{t, p, i}  \tag{5}\\
& \text { inbound }_{t, p, i}=\sum_{t^{\prime}} x_{t^{\prime}, p, i}+\sum_{t^{\prime}} \text { pur }_{t^{\prime}, p, i}+\sum_{p^{\prime}} s_{t, p^{\prime}, p, i} \\
& +\sum_{i^{\prime}, j} r p_{t, p, i, i^{\prime}, j}+\sum_{j} r_{t, p, j, i}+P O_{t, p, i}+W I P_{t, p, i} \\
& \text { outbound }_{t, p, i}=\sum_{j} B_{t, p, i, j} x_{t, p, j}+\sum_{p^{\prime}} s_{t, p, p^{\prime}, i}+\sum_{i^{\prime}, j} r p_{t, p, i^{\prime}, i, j} \\
& +\sum_{j} r_{t, p, i, j}+z_{t, p, i} \\
& \text { for }(t, p, i) \in P \\
& m_{t, p, i}=m_{t-1, p, i}-z_{t, p, i}+D_{t, p, i}  \tag{6}\\
& \text { for }(t, p, i) \in P \\
& \sum_{i^{\prime}} r p_{t, p, i, i^{\prime}, j} \leq B_{t, p, i, j} x_{t, p, j}  \tag{7}\\
& \text { for }(t, p, i, j) \in B O M
\end{align*}
$$

Indices:

- $i, i^{\prime}, j$ : items (raw material, sub-assembly or end product).
- p: plant.
- $t$ : period.

Decision variables:

- $x_{t, p, i}, p r_{t, p, i}, z_{t, p, i}$ : the production/purchase/deliver amount of item $i$ in plant $p$ at period $t$.
- $s_{t, p, p^{\prime}, t}$ : the transit amount of item $i$ from plant $p$ to plant $p^{\prime}$ at period $t$.
- $r_{t, p, i, j}$ : the amount that item $i$ replace item $j$ in plant $p$ at period $t$.
- $r p_{t, p, i, i i^{\prime}, j}$ : the amount that item $i^{\prime}$ replace item $i$ to produce item $j$ in plant $p$ at period $t$.
State variables (whose value is determined by other decision variables):
- inv $_{t, p, i}$ : the inventory amount of item $i$ in plant $p$ at period $t$.
- $m_{t, p, i}$ : the delay amount of item $i$ in plant $p$ at period $t$.

All the value of decision and state variables are not less than zero.

Some important constants (note that not all constants are listed due to space limit. Every sum operation contains a constant that controls the range of indices):

- $C^{m}, C^{s}, C^{p u r}, C^{r p}, C^{r}$ : the cost of delay, transition, purchase and replacement. (Delay cost will usually dominate the objective)
- $P[t, p, i]:$ item $i$ will be produced in plant $p$ at period $t$.
- BOM $[t, p, i, j]: j$ is the parent of $i$ in plant $p$ at period $t$.
- $B_{t, p, i, j}$ : number of item $i$ 's amount that need to be consumed to producing one item $j$.
- $P O_{t, p, i}, W I P_{t, p, i}$ : the amount of purchase order (PO)/ work-in-progress (WIP) of item $i$ in plant $p$ at period $t$.
Constraints:
- Equation 5: inventory constraint. The current inventory amount equals to last period's inventory plus inbound minus outbound.
- Equation 6: delay constraint. The current delay amount equals to last period's delay amount plus delivery amount minus demand amount.
- Equation 7: replacement constraint. For all component-assembly relation, the sum of replacement amount cannot exceed the needed amount for assembly's production.


## REFERENCES

[1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu , and Xiaoqiang Zheng. 2016. TensorFlow: A System for Large-Scale Machine Learning. In 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16). USENIX Association, Savannah, GA, 265-283. https: //www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi
[2] Ulf Adams. 2018. Ryū: Fast Float-to-String Conversion. In Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (Philadelphia, PA, USA) (PLDI 2018). Association for Computing Machinery, New York, NY, USA, 270-282. https://doi.org/10.1145/3192366.3192369
[3] MOSEK ApS. 2019. MOSEK optimization suite.
[4] Stephen P. Bradley. 1977. Applied mathematical programming. Addison-Wesley Pub. Co., Reading, Mass.
[5] Anthony Brook, David Kendrick, and Alexander Meeraus. 1988. GAMS, a user's guide. ACM Signum Newsletter 23, 3-4 (1988), 10-11.
[6] Mingyuan Chen and Weimin Wang. 1997. A linear programming model for integrated steel production and distribution planning. International fournal of Operations \& Production Management (1997).
[7] Gerard Cornuejols and Reha Tütüncü. 2006. Optimization methods in finance. Vol. 5. Cambridge University Press.
[8] A.D. Dimitriadis, N. Shah, and C.C. Pantelides. 1997. RTN-based rolling horizon algorithms for medium term scheduling of multipurpose plants. Computers \& Chemical Engineering 21 (1997), S1061 - S1066.
[9] Rafael Epstein, Andres Neely, Andres Weintraub, Fernando Valenzuela, Sergio Hurtado, Guillermo Gonzalez, Alex Beiza, Mauricio Naveas, Florencio Infante, Fernando Alarcon, et al. 2012. A strategic empty container logistics optimization in a major shipping company. Interfaces 42,1 (2012), 5-16.
[10] Robert Fourer, David M Gay, and Brian W Kernighan. 1990. A modeling language for mathematical programming. Management Science 36, 5 (1990), 519-554.
[11] Gurobi Optimization, LLC. 2018. Gurobi optimizer reference manual. https: //www.gurobi.com/documentation/9.0/refman/index.html. (2018).
[12] Gurobi Optimization, LLC. 2019. Gurobi modeling examples. https://gurobi. github.io/modeling-examples/.
[13] William E Hart, Jean-Paul Watson, and David L Woodruff. 2011. Pyomo: modeling and solving mathematical programs in Python. Mathematical Programming Computation 3, 3 (2011), 219.
[14] Thorsten Koch. 2004. Rapid Mathematical Programming. Ph.D. Dissertation. Technische Universität Berlin. http://www.zib.de/Publications/abstracts/ZR-04-

58/ ZIB-Report 04-58.
[15] Johan Lofberg. 2004. YALMIP: A toolbox for modeling and optimization in MATLAB. In 2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508). IEEE, 284-289.
[16] Miles Lubin and Iain Dunning. 2015. Computing in Operations Research Using Julia. INFORMS Journal on Computing 27, 2 (2015), 238-248. https://doi.org/10. 1287/ijoc.2014.0623
[17] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Advances in Neural Information Processing Systems, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc., 8026-8037. https://proceedings.neurips.cc/paper/2019/ file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf
[18] Yves Pochet and Laurence A. Wolsey. 2010. Production Planning by Mixed Integer Programming (1st ed.). Springer Publishing Company, Incorporated.
[19] Suresh Sethi and Gerhard Sorger. 1991. A theory of rolling horizon decision making. Annals of Operations Research 29, 1 (01 Dec 1991), 387-415. https: //doi.org/10.1007/BF02283607


[^0]:    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.
    CIKM '21, November 1-5, 2021, Virtual Event, Australia.
    © 2021 Association for Computing Machinery.
    ACM ISBN 978-1-4503-8446-9/21/11...\$15.00
    https://doi.org/10.1145/XXXXXX.XXXXXX

[^1]:    ${ }^{1}$ For example, an ad-hoc solution for instantiate Expression 3's instantiation is to pre-index the edge by both head and tail node.

[^2]:    

[^3]:    ${ }^{3}$ Note that we cannot deploy more efficient solvers like CPLEX on Huawei's online enterprise environment due to export restriction of the US (so as advanced AMSs like AMPL). However, all solvers always deliver exact optimal solution if possible, and have similar exponential curves between problem scale and solving time. Therefore the selection of external solvers will not change the optimality and solving time ratio.

[^4]:    ${ }^{4}$ Due to export restrictions, we cannot purchase and deploy the latest version of Gurobi in Huawei's enterprise environment.

