

Graph-Based Modeling and Decomposition of Energy Infrastructures^{*}

Sungho Shin^{*} Carleton Coffrin^{**} Kaarthik Sundar^{**}
Victor M. Zavala^{*}

^{*} *University of Wisconsin-Madison, Madison, WI 53706 USA*
(e-mail: sungho.shin@wisc.edu; victor.zavala@wisc.edu).

^{**} *Los Alamos National Laboratory, Los Alamos, NM 87545 USA*
(e-mail: cjc@lanl.gov; kaarthik@lanl.gov)

Abstract: Nonlinear optimization problems are found at the heart of real-time operations of critical infrastructures. These problems are computationally challenging because they embed complex physical models that exhibit space-time dynamics. We propose modeling these problems as graph-structured optimization problems, and illustrate how their structure can be exploited at the modeling level (for parallelizing function/derivative computations) and at the solver level (for parallelizing linear algebra operations). Specifically, we present a restricted additive Schwarz scheme that enables flexible decomposition of complex graph structures within an interior-point algorithm. The proposed approach is implemented as a general-purpose nonlinear programming solver that we call `MadNLP.jl`; this Julia-based solver is interfaced to the graph-based modeling package `Plasmo.jl`. The efficiency of this framework is demonstrated via problems arising in transient gas network optimization and multi-period AC optimal power flow. We show that our framework accelerates the solution (compared to off-the-shelf tools) by over 300%; specifically, solution times are reduced from 72.36 sec to 23.84 sec for the gas problem and from 515.81 sec to 149.45 sec for the power flow problem.

Keywords: Nonlinear Optimization, Decomposition, Graphs, Energy Systems

1. INTRODUCTION

Real-time operation of modern energy infrastructures requires solving large-scale nonlinear programs (NLPs). Application examples include transient gas network optimization (Sundar and Zlotnik, 2018) and multi-period optimal power flow problems (Geth et al., 2020; Kim and Anitescu, 2020). Achieving real-time solutions for these problems is challenging, as they embed complex physical models that require space-time discretization. NLPs arising in this context can easily reach millions of variables and constraints and defy the scope of off-the-shelf solvers. Specifically, scalability bottlenecks are often encountered at the modeling level (function and derivative computations) and at the solver level (computation of the search step).

Large-scale NLPs arising in energy infrastructures have the key characteristic that they exhibit *sparse graphs* structures; we refer to such problems as *graph-structured optimization problems* (Jalving et al., 2019; Shin et al., 2020b; Jalving et al., 2020). Graph-structured problems can be conveniently modeled using specialized modeling platforms such as `Plasmo.jl` (Jalving et al., 2019, 2020) and solved using structure-exploiting optimization solvers such as PIPS-NLP (Chiang et al., 2014). `Plasmo.jl` is a graph-based modeling platform that enables the modular construction and analysis of highly complex models; this platform also leverages the algebraic modeling capabilities

of `JuMP.jl` (Dunning et al., 2017) and facilitates access to infrastructure modeling tools such as `GasModels.jl` and `PowerModels.jl` (Bent et al., 2020; Coffrin et al., 2018).

Another key benefit of `Plasmo.jl` is that it can communicate model structures to solvers, and this facilitates the implementation of different decomposition strategies, notably the alternating direction method of multipliers (Boyd et al., 2011), overlapping Schwarz (Shin et al., 2020a), and parallel interior-point (IP) methods (Chiang et al., 2014; Rodriguez et al., 2020).

In this paper, we present a new and flexible decomposition framework for graph-structured optimization problems. Our framework uses a restricted additive Schwarz (RAS) decomposition scheme implemented within a filter line-search IP method (Wächter and Biegler, 2006). We present a *Julia*-based implementation of this approach, which we call `MadNLP.jl` (<https://github.com/zavalab/MadNLP.jl>). We use our framework to experiment with different decomposition strategies that exploit parallelism at the modeling and at the solver level. Specifically, we consider a scheme that parallelizes function and derivative computations by exploiting the modular structure of `Plasmo.jl`. We also consider a scheme that uses RAS (Cai and Sarkis, 1999) for parallelizing step computations. RAS has been widely used for the solution of large linear algebra systems arising from discretized partial differential equations (PDEs) (Balay et al., 2019), but we have recently shown that it can also be applied to solve general linear systems arising in graph-structured optimization (Shin et al.,

^{*} We acknowledge support from the Grainger Wisconsin Distinguished Graduate Fellowship.

2020b; Gerstner et al., 2016). Our computational results indicate that our framework can accelerate computations by up to 300% (compared to off-the-shelf tools).

The paper is organized as follows: In Section 2, we define the graph-structured problem of interest and discuss how its structure can be exploited at the modeling level. In Section 3, we discuss parallel decomposition schemes. In Section 4, we apply these schemes to transient gas network optimization and multi-period AC optimal power flow problems. Section 5 presents concluding remarks.

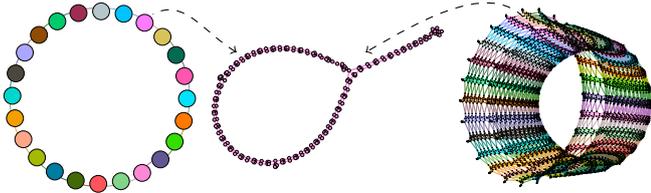


Fig. 1. Alternative graph representations of a transient gas network problem. Left: graph G in which nodes represent times and the spatial (network) structure is embedded within each node. Right: graph \tilde{G} in which each node represents a variable or a constraint. Middle: a slice of \tilde{G} that corresponds to the structure contained in each node of G .

2. GRAPH-BASED MODELING

Optimization problems arising in energy infrastructures can be expressed as a graph-structured optimization problem of the form:

$$\min_{\{x_i\}_{i \in V}} \sum_{i \in V} f_i(x_i) \quad (1a)$$

$$\text{s.t. } c_i^I(x_i) = 0, \quad i \in V, \quad (\lambda_i^I) \quad (1b)$$

$$c_i^L(\{x_j\}_{j \in N_G[i]}) = 0, \quad i \in V, \quad (\lambda_i^L) \quad (1c)$$

$$x_i \geq 0, \quad i \in V \quad (z_i) \quad (1d)$$

Here, the undirected graph $G = (V, E)$ is an ordered pair of the nonempty, strictly ordered node set V and the edge set $E \subseteq \{\{i, j\} \subseteq V : i \neq j\}$; $N_G[i] := \{j \in V : \{i, j\} \in E\} \cup \{i\}$ denotes the neighborhood of $i \in V$ on G . For each $i \in V$, $x_i \in \mathbb{R}^{n_i}$ is the decision variable; $f_i(\cdot)$ is the objective function; $c_i^I(\cdot)$ is the inner equality constraint function; $c_i^L(\cdot)$ is the linking constraint function; $\lambda_i^I \in \mathbb{R}^{m_i^I}$ is the dual variable associated with (1b); $\lambda_i^L \in \mathbb{R}^{m_i^L}$ is the dual variable associated with (1c); $z_i \in \mathbb{R}^{n_i}$ is the dual variable associated with (1d). General inequality constraints can be handled by introducing slack variables.

In the context of energy infrastructures, the graph G may intuitively be used as an abstraction of the space-time structure of the problem. Specifically, each node $i \in V$ is a component located at a particular spatial location and at a particular time point (e.g., a subset of buses, generators, storage facilities, and electric lines in power networks or a subset of junctions, compressors, and pipelines in gas networks). The link constraints (1c) may represent spatial connections (e.g., interconnecting electric lines in power networks or pipelines in gas networks) or temporal connections. Furthermore, we will show that the graph G

can also be used as a general abstraction in which each node $i \in V$ represents an individual variable or constraint of the problem (the graph encodes the general sparsity structure of the problem). The ability to represent the same optimization problem in different forms provides flexibility to identify efficient decomposition strategies.

We define the short-hand notation for (1):

$$\min_{\mathbf{x}} \mathbf{f}(\mathbf{x}) \quad (2a)$$

$$\text{s.t. } \mathbf{c}(\mathbf{x}) = 0 \quad (\boldsymbol{\lambda}) \quad (2b)$$

$$\mathbf{x} \geq 0 \quad (\mathbf{z}), \quad (2c)$$

where $\mathbf{x} := \{x_i\}_{i \in V}$, $\boldsymbol{\lambda} = \{[\lambda_i^I; \lambda_i^L]\}_{i \in V}$, $\mathbf{f}(\cdot) := \sum_{i \in V} f_i(\cdot)$, $\mathbf{c}(\cdot) := \{[c_i^I(\cdot); c_i^L(\cdot)]\}_{i \in V}$, $\mathbf{n} = \sum_{i \in V} n_i$, and $\mathbf{m} = \sum_{i \in V} m_i^I + m_i^L$. Here, n_i and m_i are the primal-dual variable dimensions, and $\{(\cdot)\}_{i \in V}$ denotes the vector concatenation. We use boldface symbols to denote quantities associated with multiple nodes.

Each node $i \in V$ may contain more than or less than one variable and constraint ($n_i, m_i \in \mathbb{Z}_{\geq 0}$). Thus, the *problem graph* $G = (V, E)$ may be different from the *primal-dual coupling graph* $\tilde{G} = (\tilde{V}, \tilde{E})$, where $\tilde{V} := \mathbb{Z}_{[1, \mathbf{n} + \mathbf{m}]}$, $\tilde{E} := \{\{i, j\} : \nabla_{(\mathbf{x}, \boldsymbol{\lambda})(\mathbf{x}, \boldsymbol{\lambda})}^2 \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{z})[i, j] \neq 0\}$. Here, $\mathcal{L}(\cdot)$ is the Lagrangian of (2), and we use syntax $\mathbb{Z}_{[a, b]} := \{a, a + 1, \dots, b\}$. We can observe that a node i in V corresponds to a set of nodes $U_i \subseteq \mathbb{Z}_{[1, \mathbf{n} + \mathbf{m}]}$ in \tilde{V} , which contains multiple variables and constraints. Example graphs G and \tilde{G} for a transient gas network problem are depicted in Figure 1 (a detailed problem formulation can be found in Section 4.1). Graph G contains 24 nodes, each corresponding to a time point in a prediction horizon. Periodicity (over a 24 hours period) is enforced as constraints (this periodicity creates the cycle shape of G). In this graph, each node embeds the spatial structure of the problem (network and pipelines). Graph \tilde{G} unfolds the temporal and spatial structure and shows the interconnectivity between all variables and constraints in the problem.

3. GRAPH-BASED DECOMPOSITION

We proceed to describe our IP solver `MadNLP.jl`, comprising a new **Schwarz** decomposition scheme that exploits graph structures within an IP method, and we describe its interface to `Plasmo.jl`.

3.1 Interior-Point Method

The IP method implemented in `MadNLP.jl` finds the solution of (2) by solving a sequence of *barrier subproblems*:

$$\min \varphi(\mathbf{x}) := \mathbf{f}(\mathbf{x}) - \mu \mathbf{e}^T \log(\mathbf{x}) \quad \text{s.t. } \mathbf{c}(\mathbf{x}) = 0. \quad (3)$$

with a decreasing sequence for parameter μ . The KKT conditions for (3) give the nonlinear equations:

$$\nabla \mathbf{f}(\mathbf{x}) + A^T \boldsymbol{\lambda} - \mathbf{z} = 0; \quad \mathbf{c}(\mathbf{x}) = 0; \quad \mathbf{X} \mathbf{Z} \mathbf{e} - \mu \mathbf{e} = 0, \quad (4)$$

where $A := \nabla \mathbf{c}(\mathbf{x})$, $\mathbf{X} := \text{diag}(\mathbf{x})$, and $\mathbf{Z} := \text{diag}(\mathbf{z})$. A solution of KKT system (4) is obtained by computing primal-dual Newton steps \mathbf{d}^* from:

$$\underbrace{\begin{bmatrix} W + \Sigma + \delta_w I & A^T \\ A & -\delta_c I \end{bmatrix}}_M \underbrace{\begin{bmatrix} d^x \\ d^\lambda \end{bmatrix}}_d = - \underbrace{\begin{bmatrix} \nabla_{\mathbf{x}} \varphi(\mathbf{x}) + A^T \boldsymbol{\lambda} \\ \mathbf{c}(\mathbf{x}) \end{bmatrix}}_p, \quad (5)$$

where $W := \nabla_{\mathbf{x}\mathbf{x}}^2 \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{z})$, $\Sigma := \mathbf{X}^{-1} \mathbf{Z}$, and $\delta_w, \delta_c > 0$ are regularization parameters. The step \mathbf{d}^* computed from (5) is safeguarded by a line-search filter procedure to induce global convergence (Wächter and Biegler, 2006). Typically, the solution of the linear system (5) is the most computationally intensive step in the IP method. This system is typically solved using direct linear solvers that are based on LDL^\top factorizations (e.g., as those implemented in HSL routines (HSL, 2007)). Decomposition strategies based on Schur complements (Chiang et al., 2014) and iterative strategies (Curtis et al., 2012; Rodriguez et al., 2020) have also been proposed. The solution of (5) based on a direct block LDL^\top factorization reveals the inertia (the number of positive, zero, negative eigenvalues) of \mathbf{M} . This inertia information is crucial in determining the acceptability of the computed step and in triggering the regularization of the linear system. However, inertia is not available when using iterative solution algorithms (as that proposed in this work). In MadNLP.jl, we use an inertia-free regularization strategy to determine the acceptability of the step (Chiang and Zavala, 2016). This method performs a simple negative curvature test to trigger regularization.

3.2 Restricted Additive Schwarz (RAS)

We propose a solution strategy for (5) based on an RAS scheme. We define some concepts and quantities that help explain our algorithm. Consider a partition $\{V_k\}_{k=1}^K$ of V ; we call V_k *non-overlapping subdomains*. This partition can be obtained by applying a graph partitioning scheme to G . We then construct a family of *overlapping subdomains* $\{V_k^\omega\}_{k=1}^K$ (these are constructed by *expanding* V_k). The expansion procedure is performed by progressively incorporating adjacent nodes (Shin et al., 2020b) (the *size of overlap* ω represents the expansion level). We observe that:

$$V_k \subseteq V_k^\omega \subseteq V, \quad k = 1, \dots, K; \quad \bigsqcup_{k=1}^K V_k = \bigcup_{k=1}^K V_k^\omega = V,$$

where \sqcup denotes disjoint union. With $\{V_k\}_{k=1}^K$ and $\{V_k^\omega\}_{k=1}^K$, we define the corresponding index sets in the space of primal-dual variables in \mathbb{R}^{n+m} as follows:

$$W_k := \bigsqcup_{i \in V_k} U_i; \quad W_k^\omega := \bigsqcup_{i \in V_k^\omega} U_i, \quad k = 1, \dots, K,$$

where $U_i \subseteq W := \mathbb{Z}_{[1, n+m]}$ is the index set of $[x_i; \lambda_i]$ in $[\mathbf{x}; \boldsymbol{\lambda}]$. We now observe that:

$$W_k \subseteq W_k^\omega \subseteq W, \quad k = 1, \dots, K; \quad \bigsqcup_{k=1}^K W_k = \bigcup_{k=1}^K W_k^\omega = W.$$

We state the RAS scheme for solving (5) as:

$$\mathbf{d}^{(\ell+1)} = \mathbf{d}^{(\ell)} + \underbrace{\left(\sum_{k=1}^K \tilde{R}_k \mathbf{M}_k^{-1} R_k^\top \right)}_{P^{-1}} \mathbf{r}^{(\ell)}, \quad \ell = 0, 1, \dots \quad (6)$$

Here, ℓ is the RAS iteration counter, $\mathbf{r}^{(\ell)} := \mathbf{p} - \mathbf{M} \mathbf{d}^{(\ell)}$ is the residual; $\mathbf{M}_k := \mathbf{M}[W_k^\omega, W_k^\omega]$; $R_k := \{e_i^\top\}_{i \in W_k}$; $\tilde{R}_k := \{\tilde{e}_i^k\}_{i \in W_k}^\top$, where e_i is the i -th standard basis of \mathbb{R}^{n+m} , and $\tilde{e}_i^k = e_i$ if $i \in W_k$ and 0 otherwise.

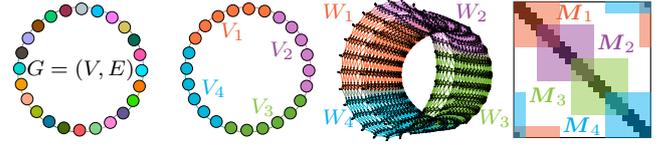


Fig. 2. Graph-based decomposition in MadNLP.jl applied to a transient gas network problem.

The RAS scheme (6) involves the following steps. We first obtain the residual at the current step ℓ ; then, for each *overlapping* subdomains $\{V_k^\omega\}_{k=1}^K$, the associated residual is extracted as $R_k^\top \mathbf{r}^{(\ell)}$. The k -th subsystem is then solved by applying \mathbf{M}_k^{-1} . In MadNLP.jl, a factorization of \mathbf{M}_k is computed with a direct solver and stored, so that the system can be repeatedly solved whenever the new right-hand-side is given. Subsequently, the solution $\mathbf{M}_k^{-1} R_k^\top \mathbf{r}^{(\ell)}$ for the k -th overlapping subdomain is *restricted* to the non-overlapping subdomain V_k and then mapped back to the full-space by applying \tilde{R}_k (the indices associated with $V \setminus V_k$ are set to zero in this step). Key defining features of RAS are the concepts of *overlap* and *restriction*. The overlap allows the dampening of the adverse effect of the truncated domain, and the restriction procedure discards the part of the solution where the adverse truncation effect is strong.

It has been empirically observed that the convergence of the RAS algorithm improves as the size of overlap ω increases and as the conditioning of \mathbf{M} improves (Cai and Saad, 1996). For positive definite \mathbf{M} , an exponential relationship between the convergence rate and the size of overlap has been established (Shin et al., 2020b). When $\omega = 0$, the RAS scheme reduces to a block-Jacobi scheme (decentralized) while, when ω is maximal ($\mathbf{M}_k = \mathbf{M}$), the RAS scheme becomes a direct solution method (centralized). In this sense, RAS provides a bridge between fully decentralized and fully centralized schemes (thus providing flexibility). In the Schwarz submodule of MadNLP.jl, ω is set automatically based on the relative size of V_k , and adaptively adjusted whenever a convergence issue occurs. Algorithm (6) uses a simple static iteration (also called a Richardson iteration), but more sophisticated iterative methods (e.g., the generalized mean residual (GMRES) method) can also be used (by treating P as a preconditioner). Both Richardson and GMRES iterators are implemented in MadNLP.jl.

The construction of partitions for the RAS scheme is illustrated in Figure 2; these partitions correspond to the transient gas network example of Figure 1. Here, by partitioning the problem graph $G = (V, E)$ (first subfigure), the node set V is divided into 4 subdomains V_1, V_2, V_3, V_4 (second subfigure). These subdomains are associated with the primal-dual index sets W_1, W_2, W_3, W_4 (third subfigure). After applying expansions, the associated blocks $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3, \mathbf{M}_4$ of \mathbf{M} are identified (last subfigure); these blocks are used by the RAS scheme (6).

3.3 Implementation in MadNLP.jl

The abstraction layers within MadNLP.jl are shown in Figure 3. The problem is modeled as an OptiGraph object (the core modeling object of PlasmO.jl). The OptiGraph

is interfaced with multiple JuMP.jl models. The JuMP.jl model objects provide a set of local function oracles (objective, objective gradient, constraint, constraint Jacobian, and Lagrangian hessian). The PlasmO.jl-MadNLP.jl interface collects these local function oracles and creates a set of oracles for the full problem, where the local oracles are evaluated in parallel.

The Solver object of MadNLP.jl is created from the OptiGraph object of PlasmO.jl. The Solver object of MadNLP.jl uses a line-search filter IP method (Wächter and Biegler, 2006) to solve the problem. The step computation is performed by the linear solver specified by the user. The linear solver can be specified either as a direct solver or as the RAS solver. When the RAS solver is chosen, multiple subproblem solver objects are created by using standard direct solvers (e.g., by using Ma57 of HSL routines). These subproblem solvers are used for factorization and backsolve for M_k blocks. The RAS scheme (6) exploits multi-thread parallelism available in Julia. After termination of the IP solution procedure, the primal-dual solutions are sent back to the OptiGraph object and Model objects from JuMP.jl so that the user can query the solution via the interface provided by PlasmO.jl and JuMP.jl. See Figure 3 for a comparison with a conventional implementation.

4. CASE STUDIES

The proposed computational framework was tested using transient gas network and multi-period AC power flow problems. In this section, we present the problem statements followed by numerical results and discussion.

4.1 Transient Gas Network

We consider a transient gas network problem (Sundar and Zlotnik, 2018) of the form:

$$\min_{\substack{\rho, \varphi^a, \varphi^-, \\ \alpha, s, d \in \mathbb{R}}} \sum_{t \in \mathcal{T}} \left(\sum_{(i,j) \in \mathcal{C}} \gamma P_{ijt}^a + \sum_{i \in \mathcal{R}} c_{it} s_{it} - \sum_{i \in \mathcal{D}} c_{it} d_{it} \right) \quad (7a)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{N}(i)} f_{ijt}^a = \sum_{j \in \mathcal{R}(i)} s_{jt} - \sum_{j \in \mathcal{D}(i)} d_{jt}, i \in \mathcal{N}, t \in \mathcal{T} \quad (7b)$$

$$\rho_{it}^{\min} \leq \rho_{it} \leq \rho_{it}^{\max}, i \in \mathcal{N}, t \in \mathcal{T} \quad (7c)$$

$$\rho_{it}^2 - \rho_{jt}^2 = -\frac{\lambda L}{D} \varphi_{ijt}^a |\varphi_{ijt}^a|, (i, j) \in \mathcal{P}, t \in \mathcal{T} \quad (7d)$$

$$\hat{L}(\dot{\rho}_{jt} + \dot{\rho}_{it}) = -4\varphi_{ijt}^-, (i, j) \in \mathcal{P}, t \in \mathcal{T} \quad (7e)$$

$$f_{ijt}^a (\rho_{it} - \rho_{jt}) \leq 0, (i, j) \in \mathcal{C}, t \in \mathcal{T} \quad (7f)$$

$$\rho_{jt} = \alpha_{ijt} \rho_{it}, (i, j) \in \mathcal{C}, t \in \mathcal{T} \quad (7g)$$

$$P_{ijt}^a \leq P_{ijt}^{\max}, (i, j) \in \mathcal{C}, t \in \mathcal{T} \quad (7h)$$

$$-f_{ij}^a \leq f_{ijt}^a \leq f_{ij}^a, (i, j) \in \mathcal{C}, t \in \mathcal{T}, \quad (7i)$$

where $\dot{\rho}_{it} = \frac{\rho_{it} - \rho_{it-1}}{\Delta t}$, $P_{ijt}^a = W_a A_{ij}$, and $f_{ijt}^a = A_{ij} \varphi_{ijt}^a$. Here, \mathcal{N} is the set of junctions; \mathcal{P} is the set of pipelines; \mathcal{C} is the set of compressors; \mathcal{R} is the set of receipts; \mathcal{D} is the set of demands; $\mathcal{R}(i)$ is the set of receipts at junction $i \in \mathcal{N}$; $\mathcal{D}(i)$ is the set of demands at junction $i \in \mathcal{N}$; \mathcal{T} is the time index set; ρ is the density; φ^a is the average mass flux; φ^- is the negative mass flux; α is the compression ratio; s is the supply; d is the

demand; $\dot{\rho}$ is the time derivative of density; P^a is the power consumption of compressor; f is the mass flow; c is the gas price; γ is the economic factor; $\lambda, \hat{L}, L, D, A, \Delta t$, and W_a are physical parameters. To implicitly enforce the periodicity, we let $\rho_{i0} = \rho_{iT}$, where T is the end time index. The gas network under study consists of 2 compressors, 6 junctions (35 junctions after discretization), 4 pipelines (32 pipelines after discretization), 1 receiving points and 5 transfer points (which work either as receipt or delivery). The model is constructed using GasModels.jl (Bent et al., 2020).

4.2 Multi-Period AC Power Flow

We consider a multi-period AC power flow problem with storage (Geth et al., 2020) of the form:

$$\min_{\substack{v, s, s^g, s^s \in \mathbb{C} \\ e, sc, sd, sqc \in \mathbb{R}}} \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{G}} c_{kt}^0 + c_{kt}^1 \Re(s_{kt}^g) + c_{kt}^2 \Re(s_{kt}^g)^2 \quad (8a)$$

$$\text{s.t.} \quad v_i^L \leq |v_{it}| \leq v_i^U, i \in \mathcal{N}, t \in \mathcal{T} \quad (8b)$$

$$\sum_{k \in \mathcal{G}_i} s_{kt}^g - \sum_{k \in \mathcal{L}_i} s_{kt}^d + \sum_{k \in \mathcal{S}_i} s_{kt}^s = \sum_{j \in \mathcal{N}_G[i]} s_{ijt}, i \in \mathcal{N}, t \in \mathcal{T} \quad (8c)$$

$$s_{ijt} = (Y_{ij} + Y_{ij}^c)^* \frac{|v_{it}|^2}{|T_{ij}|^2} - Y_{ij}^* \frac{v_{it} v_{jt}^*}{T_{ij}}, \quad (i, j) \in \mathcal{E}, t \in \mathcal{T} \quad (8d)$$

$$s_{ijt} = (Y_{ij} + Y_{ji}^c)^* |v_{jt}|^2 - Y_{ij}^* \frac{v_{it} v_{jt}^*}{T_{ij}^*}, \quad (i, j) \in \mathcal{E}^R, t \in \mathcal{T} \quad (8e)$$

$$|s_{ijt}| \leq s_{ij}^U, (i, j) \in \mathcal{E} \cup \mathcal{E}^R, t \in \mathcal{T} \quad (8f)$$

$$\theta_{ij}^{\Delta L} \leq \angle(v_{it} v_{jt}^*) \leq \theta_{ij}^{\Delta U}, (i, j) \in \mathcal{E}, t \in \mathcal{T} \quad (8g)$$

$$s_k^{gL} \leq s_{kt}^g \leq s_k^{gU}, k \in \mathcal{G}, t \in \mathcal{T} \quad (8h)$$

$$e_{kt} - e_{kt-1} = (\eta^c sc_t - sd_t / \eta^d) \Delta t, k \in \mathcal{S}, t \in \mathcal{T} \setminus \{T\} \quad (8i)$$

$$s_{kt}^s + (sc_{kt} - sd_{kt}) = \sqrt{-1} sqc_{kt} + s_k^{\text{loss}}, k \in \mathcal{S}, t \in \mathcal{T} \quad (8j)$$

$$|s_{kt}^s| \leq s_k^u, 0 \leq e_{kt} \leq e_k^u, k \in \mathcal{S}, t \in \mathcal{T} \quad (8k)$$

$$0 \leq sc_{kt} \leq sc_k^u, 0 \leq sd_{kt} \leq sd_k^u, k \in \mathcal{S}, t \in \mathcal{T}, \quad (8l)$$

Here, \mathcal{G} is the set of generators; \mathcal{N} is the set of buses; \mathcal{E} is the set of (directed) branches; \mathcal{E}^R is the set of branches with inverted directions; \mathcal{S} is the set of storage; \mathcal{T} is the time index set; $v \in \mathbb{C}$ is the voltage; $e \in \mathbb{R}$ is the state of charge; $s \in \mathbb{C}$ is the power flow; $s^g \in \mathbb{C}$ is the power generation; $s^s \in \mathbb{C}$ is the complex power injected by the storage; $sc \in \mathbb{R}$ is the charging rate; $sd \in \mathbb{R}$ is the discharging rate; $sqc \in \mathbb{R}$ is the reactive power slack; $c^0, c^1, c^2 \in \mathbb{R}$ are the generation costs; $s^d \in \mathbb{C}$ is the power demand; Y is the admittance; T is the branch complex transformation parameter; η is the charging efficiency; s^{loss} is the storage energy loss; Δt is the time interval. Note that (8) can be reformulated as an NLP with real variables by separately treating the real and imaginary part of the variables and equations (a polar formulation is used here). The power network under study is a variant of IEEE 14 bus test system; this comprises 14 buses, 5 generators, 1 storage, 1 shunt, and 20 branches. The detailed model is constructed with PowerModels.jl (Coffrin et al., 2018).

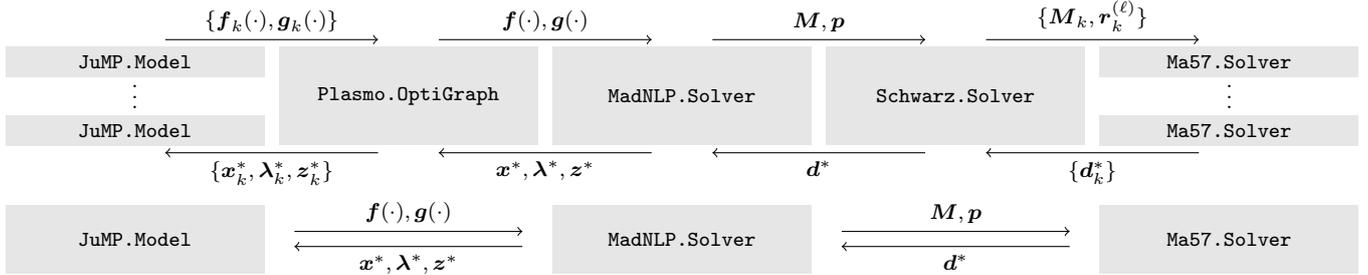


Fig. 3. Schematics of graph-based modeling and solution (top) and conventional modeling and solution (bottom).

4.3 Results and Discussion

We compare the proposed method (MadNLP.jl interfaced with PlasmO.jl and Schwarz), with the conventional method (MadNLP.jl interfaced with serial/parallel direct solvers Ma57 or MKL-Pardiso along with non-graph based algebraic modeling language JuMP.jl). The conventional methods are referred to as JuMP-Ma57 and JuMP-PardisoMKL, and the proposed method is referred to as PlasmO-Schwarz/Ma57. Furthermore, the mix of proposed/conventional approaches (JuMP-Schwarz/Ma57, PlasmO-Ma57, and PlasmO-PardisoMKL) is also tested together. For JuMP-Schwarz/Ma57, the graph partitioning tool METIS was used to partition the primal-dual coupling graph \tilde{G} directly. A Richardson scheme was used as an iterator for the RAS scheme. The study was performed by solving the gas (7) and power (8) problems while varying the size of the problems (by increasing the length of the prediction horizon). The code was run on a server computer equipped with 2 CPUs of Intel Xeon CPU E5-2698 v4 running 2.20GHz (20 core for each), and 20 threads are used for the computation. Code to reproduce the results can be found in <https://github.com/zavalab/JuliaBox/tree/master/AdchemCaseStudy>

For both problems, we found that the graph-based approach can significantly accelerate the solution (see Figure 4). In particular, comparing JuMP-Ma57 and PlasmO-Schwarz/Ma57, PlasmO-Schwarz/Ma57 becomes faster than JuMP-Ma57 when the prediction horizon is 3 days or more. Function evaluations are always faster in PlasmO.jl compared to JuMP.jl because the computational savings from function evaluations directly reduce the total solution time (parallelizing the function evaluation itself has no impact on the other part of the algorithm). On the other hand, one can see that the speed-up from parallel linear algebra is only observed when the problem size is sufficiently large (3 days in the gas network and 60 days in the power network). This is because the reduction in the problem size also reduces the overlap size. In our implementation, we set the size of overlap using the relative size of the block (the size of the overlap is reduced if the overall problem size is reduced). As a result, the RAS scheme (6) may become slow, and the number of required factorization/backsolve steps increases. This indicates that the use of RAS is beneficial only when the problem size is sufficiently large. For the gas problems, the acceleration of linear algebra computations was more pronounced. In contrast, for the power problems, the acceleration of func-

tion evaluations was more pronounced. This is because the AC power flow formulation has a large number of nonlinear expressions. By comparing the linear solver time for JuMP-Schwarz/Ma57 and PlasmO-Schwarz/Ma57, we see the advantage of using a graph-based modeling language for obtaining the partitions $\{W_k\}_{k=1}^K$. We recall that for JuMP-Schwarz/Ma57, the Metis graph partitioning routine is directly applied to \tilde{G} while PlasmO-Schwarz/Ma57 uses the user-provided problem graph G . One can observe that, in general, the linear solver time is shorter for PlasmONLP-Schwarz/Ma57. This indicates that the user-provided graph information can be leveraged for obtaining high-quality partitions.

5. CONCLUSIONS AND FUTURE WORK

We have presented a graph-based modeling and decomposition framework for large-scale nonlinear programs arising in energy infrastructures. Here, we introduce a new decomposition paradigm for linear algebra systems within an interior-point method: a restricted additive Schwarz (RAS) scheme. We implement this framework in the Julia package MadNLP.jl and show that the RAS approach accelerates computations (compared to off-the-shelf approaches) by up to 300%. This work focused on applying RAS to conduct temporal decomposition; applying RAS as a spatial decomposition scheme is a future direction of interest. A surprising finding was that the RAS scheme is effective at handling instances with a large number of active inequality constraints. We are interested in determining the reasons for this by investigating the convergence properties of the RAS scheme within an interior-point context.

ACKNOWLEDGEMENTS

We acknowledge Andreas Wächter for helpful comments and Jordan Jalving for implementing the requested features in PlasmO.jl.

REFERENCES

- Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W., et al. (2019). PETSc users manual.
- Bent, R., Sundar, K., and Fobes, D. (2020). GasModels.jl. <https://github.com/lanl-ansi/GasModels.jl>.
- Boyd, S., Parikh, N., and Chu, E. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc.

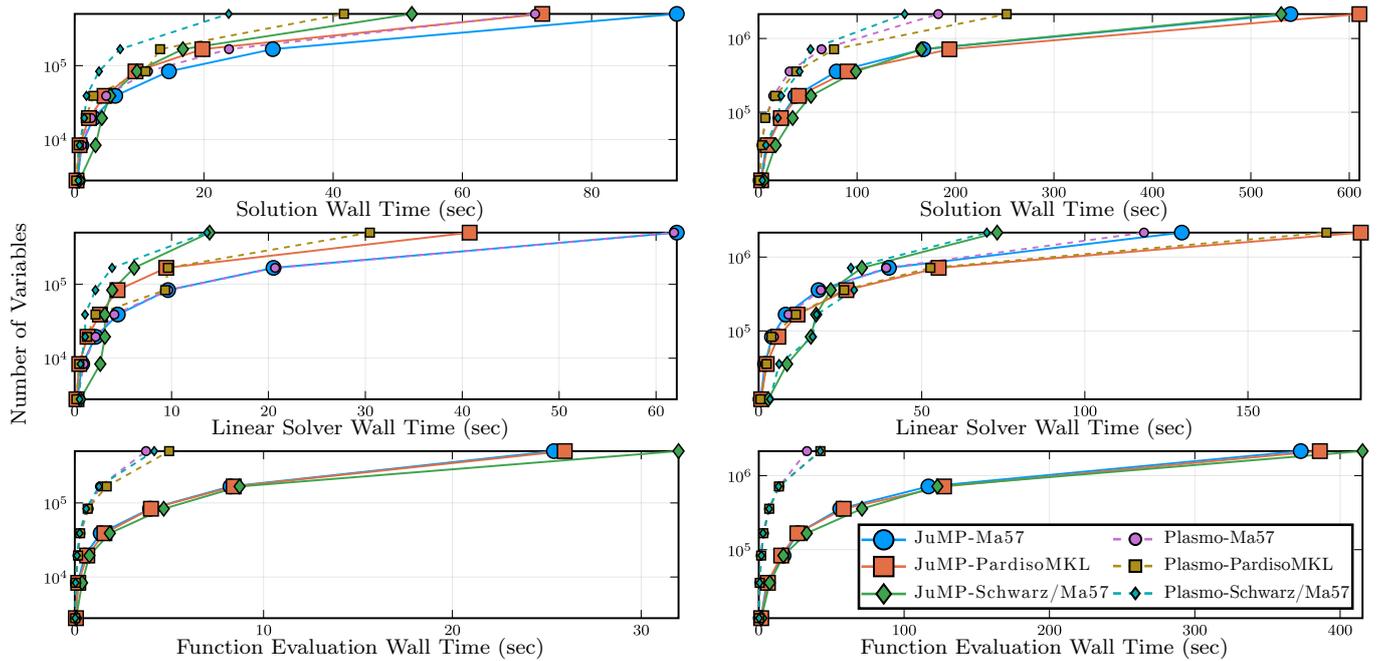


Fig. 4. Solution time (top), linear solver time (middle), function evaluation time (bottom) for transient gas network (left) and multi-period AC optimal power flow (right) problems.

- Cai, X.C. and Saad, Y. (1996). Overlapping domain decomposition algorithms for general sparse matrices. *Numerical linear algebra with applications*, 3(3), 221–237.
- Cai, X.C. and Sarkis, M. (1999). A restricted additive schwarz preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing*, 21(2), 792–797.
- Chiang, N.Y. and Zavala, V.M. (2016). An inertia-free filter line-search algorithm for large-scale nonlinear programming. *Computational Optimization and Applications*, 64(2), 327–354.
- Chiang, N., Petra, C.G., and Zavala, V.M. (2014). Structured nonconvex optimization of large-scale energy systems using PIPS-NLP. In *2014 Power Systems Computation Conference*, 1–7. IEEE.
- Coffrin, C., Bent, R., Sundar, K., Ng, Y., and Lubin, M. (2018). PowerModels.jl: An open-source framework for exploring power flow formulations. In *2018 Power Systems Computation Conference (PSCC)*, 1–8. doi: 10.23919/PSCC.2018.8442948.
- Curtis, F.E., Huber, J., Schenk, O., and Wächter, A. (2012). A note on the implementation of an interior-point algorithm for nonlinear optimization with inexact step computations. *Mathematical programming*, 136(1), 209–227.
- Dunning, I., Huchette, J., and Lubin, M. (2017). JuMP: A modeling language for mathematical optimization. *SIAM Review*, 59(2), 295–320.
- Gerstner, P., Schick, M., Heuveline, V., Meyer-Hübner, N., Suriyah, M., Leibfried, T., Slednev, V., Fichtner, W., and Bertsch, V.V. (2016). A domain decomposition approach for solving dynamic optimal power flow problems in parallel with application to the german transmission grid. *Preprint Series of the Engineering Mathematics and Computing Lab*, (1).
- Geth, F., Coffrin, C., and Fobes, D.M. (2020). A flexible storage model for power network optimization. *arXiv preprint arXiv:2004.14768*.
- HSL, A. (2007). collection of fortran codes for large-scale scientific computation. See <http://www.hsl.rl.ac.uk>.
- Jalving, J., Cao, Y., and Zavala, V.M. (2019). Graph-based modeling and simulation of complex systems. *Computers & Chemical Engineering*, 125, 134–154.
- Jalving, J., Shin, S., and Zavala, V.M. (2020). A graph-based modeling abstraction for optimization: Concepts and implementation in plasmO. jl. *arXiv preprint arXiv:2006.05378*.
- Kim, Y. and Anitescu, M. (2020). A real-time optimization with warm-start of multiperiod ac optimal power flows. *Electric Power Systems Research*, 189, 106721.
- Rodriguez, J.S., Laird, C.D., and Zavala, V.M. (2020). Scalable preconditioning of block-structured linear algebra systems using adm. *Computers & Chemical Engineering*, 133, 106478.
- Shin, S., Anitescu, M., and Zavala, V.M. (2020a). Overlapping schwarz decomposition for constrained quadratic programs. *arXiv preprint arXiv:2003.07502*.
- Shin, S., Zavala, V.M., and Anitescu, M. (2020b). Decentralized schemes with overlap for solving graph-structured optimization problems. *IEEE Transactions on Control of Network Systems*.
- Sundar, K. and Zlotnik, A. (2018). State and parameter estimation for natural gas pipeline networks using transient state data. *IEEE Transactions on Control Systems Technology*, 27(5), 2110–2124.
- Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1), 25–57.