arXiv:2108.01219v1 [math.OC] 2 Aug 2021

Computing the Newton-step faster than Hessian accumulation

Akshay Srinivasan [1] **Emanuel Todorov** [2]

Abstract

Computing the Newton-step of a generic function with N decision variables takes $O(N^3)$ flops. In this paper, we show that given the computational graph of the function, this bound can be reduced to $O(m\tau^3)$, where τ, m are the width & size of a tree-decomposition of the graph. The proposed algorithm generalizes non-linear optimal-control methods based on LQR to general optimization problems and provides non-trivial gains in iteration-complexity even in cases where the Hessian is dense.

1. Introduction

Newton's method forms the basis for second-order methods in optimization. Computing the Hessian of a generic function $f : \mathbb{R}^N \to \mathbb{R}$, requires $O(N^2)$ flops; inverting this matrix requires a further $O(N^3)$ flops. This super-linear scaling in compute/memory requirements is a major obstacle in the application of such methods despite their quadratic convergence.

The iteration-complexity of second-order methods has necessitated the development of specialized algorithms for restricted classes of problems. In particular, *Differential Dynamic Programming* (DDP) methods have been proposed for solving nonlinear optimal-control problems, and these methods, quite miraculously, achieve quadratic convergence despite their linear iteration-complexities (Jacobson & Mayne, 1970) (De O. Pantoja, 1988) (Wright, 1990).

DDP-methods were historically derived using *Dynamic Programming* (DP) and resemble LQR control-design. They were later, quite surprisingly, also found to be related to both, the Newton-iteration on the unconstrained problem (De O. Pantoja, 1988), and the SQP-iteration on the constrained problem (Wright, 1990). Sadly, the method is not trivially generalized to other objective functions.

Applications of DP to other domains, esp. the solution of least-squares problems and inference in graphical models, have seen great success, but such techniques are not applicable to problems in optimal-control or neural-networks, due to their inability to handle function compositions.

In this paper, we develop a method that exploits the compositional structure in a given objective function in order to automatically derive a *fast* Newton update. This is done by extending the connections of DDP to constrained & unconstrained optimization using tools from *Automatic Differentiation* (AD), and by using techinques from *Sparse Linear Algebra* (SLA) to bound iteration-complexity.

2. Problem setup

2.1. Computational graph

Let $\mathcal{G}$ be a Directed Acyclic Graph (DAG). Define,

$$\begin{aligned} \mathrm{pa}(u) &\triangleq \{v | (v,u) \in E[\mathcal{G}]\}, \quad \text{(parents)} \\ \mathrm{ch}(u) &\triangleq \{v | (u,v) \in E[\mathcal{G}]\}. \quad \text{(children)} \end{aligned} \tag{1}$$

Let every vertex $v \in V[\mathcal{G}]$ be associated with a *state* $\mathtt{X}_v \in U_v \subset \mathbb{R}^{n_v}$ for some open set U_v. Let $\mathtt{X}_A$ be the (labelled) concatenation of *state*s associated with vertices in set $A \subset V[\mathcal{G}]$.

Let the *input* nodes $\mathtt{Input} \triangleq \{u_1, u_2, \ldots, u_n\} \subset V[\mathcal{G}]$ be the parentless vertices in $\mathcal{G}$. Let the *states* of the non-input nodes be defined recursively by $\mathtt{X}_v \triangleq \Phi_v(\mathtt{X}_{\mathrm{pa}(v)})$ for some given function $\Phi_v : \prod_{z\in \mathrm{pa}(v)} U_z \to U_v$. Since $\mathcal{G}$ is a DAG, it follows that $\mathtt{X}_{V[\mathcal{G}]}$ is uniquely determined from the input state $\mathtt{X}_{\mathtt{Input}}$ and functions $\{\Phi_v\}_{v\in V[\mathcal{G}]\setminus X}$.

An objective function $f : U_{x_1} \times \ldots U_{x_n} \to \mathbb{R}$ has the *computational structure* given by the tuple $(\mathcal{G}, \{\Phi_v\}_{v\in V[\mathcal{G}]\setminus \mathtt{Input}}, \{\ell_v\}_{v\in V[\mathcal{G}]})$ if it can be written as the sum of local objectives $\ell_v : \prod_{z\in\{v\}\cup \mathrm{pa}(v)} U_z \to \mathbb{R}$ on the graph $\mathcal{G}$ in the following form (2),

$$\begin{aligned} f : (\mathtt{X}_{x_1}, \ldots, \mathtt{X}_{x_n}) &\mapsto \sum_{v\in V[\mathcal{G}]} \ell_v(\mathtt{X}_{v\cup \mathrm{pa}(v)}), \\ \mathtt{X}_v &\leftarrow \Phi_v(\mathtt{X}_{\mathrm{pa}(v)}), \quad \forall v \in V[\mathcal{G}], \mathrm{pa}(v) \neq \emptyset. \end{aligned} \tag{2}$$

[1]SonyAI, Tokyo, Japan. [2]University of Washington, Seattle, WA, USA. Correspondence to: Akshay Srinivasan <akssri@vakra.xyz>.

Proceedings of the *Beyond First Order Methods in Machine Learning Systems* workshop, ICML 2021. Copyright 2021 by the author(s).

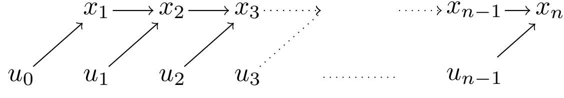

Figure 1. (Optimal control) The nodes $\{x_i\}, \{u_i\}$ represent the states and control inputs of the dynamical system.

The *assignment* operator, '$\leftarrow$', is explicitly distinguished from the equality operator, '$=$', which is taken to represent a constraint in the program. We refer to the DAG $\mathcal{G}$ as the *computational graph* of $f(\cdot)$.

NOTATION. The symbolism $\partial_u v \triangleq \frac{\partial \mathtt{X}_v}{\partial \mathtt{X}_u}\big|_{\mathtt{X}_{\mathtt{Input}}}$, will be employed for succintly denoting partial derivatives. The derivative operator *w.r.t* the (labelled) set $A = \{v_1, v_2, \dots\} \subset V[\mathcal{G}]$, will similarly be denoted by $\partial_A \triangleq [\partial_{a_1}, \partial_{a_2}, \dots]$.

2.2. Motivation.

Consider the canonical optimal-control problem,

$$\min_{u_0, u_1, \dots, u_{n-1}} \left[\mathcal{J}(u_0, \dots, u_n) \triangleq \sum_{i=0}^{n-1} \ell_i(x_i, u_i) + \ell_n(x_n) \right], \quad \forall i, x_{i+1} \leftarrow \mathtt{f}(x_i, u_i), \tag{3}$$

wherein, $\mathtt{f}(\cdot,\cdot)$ denotes the system dynamics, and $\ell_i(\cdot,\cdot)$'s denote the losses incurred. The order of computation for the objective (3) can be represented by a linear-chain as shown in Figure 1.

The computational graph Figure 1, while sparse, does not imply sparsity in the objective. Substitution of the symbol assignments in the unconstrained objective, has a cascading effect whereby variables get *transport*ed down $\mathcal{G}$,

$$\begin{aligned} \mathcal{J}(u_0, \dots, u_n) = &\ell_0(x_0, u_0) + \\ &\ell_1(\mathtt{f}(x_0, u_0), u_1) + \\ &\ell_2(\mathtt{f}(\mathtt{f}(x_0, u_0), u_1), u_2) + \dots. \end{aligned}$$

This results in a Hessian that is both fully dense and (very) badly conditioned. These two properties render both iterative and direct-factorization methods unsuitable for solving this problem.

The computational graph does, however, fully encode the sparsity of the Lagrangian of the constrained program; the constrained form being obtained by replacing '$\leftarrow$' with '$=$', and including $\{x_i\}'s$ in the domain of optimization. This results in a *fast* linear-time SQP (Sequential Quadratic Programming) iteration for the optimal-control problem but increases implementation complexity by requiring *dual* and *non-input* updates.

This presents one with a strange set of choices: slow-simple unconstrained optimization or fast-complex constrained optimization. DDP-like methods resolve this dichotomy for optimal-control problems (Jacobson & Mayne, 1970) (De O. Pantoja, 1988) (Wright, 1990). Our goal is to do the same for general problems.

3. Graphical Newton

Consider the objective function in (2), defined by the tuple $(\mathcal{G}, \{\Phi_v\}, \{\ell_v\})$. The optimization problem of interest is the following,

$$\min_{\{\mathtt{X}_v | \forall v \in \mathtt{Input}\}} \left(f \triangleq \sum_{v \in V[\mathcal{G}]} \ell_v(\mathtt{X}_{v \cup \mathrm{pa}(v)}) \right), \quad \mathtt{X}_v \leftarrow \Phi_v(\mathtt{X}_{\mathrm{pa}(v)}), \quad \forall v \in V[\mathcal{G}], \mathrm{pa}(v) \neq \emptyset. \tag{4}$$

The corresponding constrained problem is obtained by simply replacing the operator '$\leftarrow$' with '$=$' in (4) and enlarging the domain of optimization.

$$\min_{\{\mathtt{X}_v | \forall v \in V[\mathcal{G}]\}} \left(f \triangleq \sum_{v \in V[\mathcal{G}]} \ell_v(\mathtt{X}_{v \cup \mathrm{pa}(v)}) \right), \quad \mathtt{X}_v = \Phi_v(\mathtt{X}_{\mathrm{pa}(v)}), \quad \forall v \in V[\mathcal{G}], \mathrm{pa}(v) \neq \emptyset. \tag{5}$$

The constrained program (5) induces the following Lagrangian function,

$$\mathcal{L}(\mathtt{X}_{V[\mathcal{G}]}, \lambda_{V[\mathcal{G}] \setminus \mathtt{Input}}) \triangleq \sum_{v \in V[\mathcal{G}]} \ell_v(\mathtt{X}_{v \cup \mathrm{pa}(v)}) + \sum_{\substack{v \in V[\mathcal{G}], \\ \mathrm{pa}(v) \neq \emptyset}} \lambda_v^{\mathrm{T}} \mathbf{h}_v(\mathtt{X}_{v \cup \mathrm{pa}(v)}),$$

$$\text{where, } \mathbf{h}_v(\mathtt{X}_{v \cup \mathrm{pa}(v)}) \triangleq \Phi_v(\mathtt{X}_{\mathrm{pa}(v)}) - \mathtt{X}_v, \ \forall v \in V[\mathcal{G}], \mathrm{pa}(v) \neq \emptyset. \tag{6}$$

The necessary first-order conditions for optimality are given by,

$$\partial_V \mathcal{L}(\mathtt{X}_V^*, \lambda_V^*) = 0, \quad h(\mathtt{X}_V^*) = 0. \tag{7}$$

Linearization of the first-order conditions for this constrained problem (around a nominal $(\tilde{\mathtt{X}_V}, \tilde{\lambda})$) yields a system of KKT equations, whose solution yields the SQP search direction.

$$\begin{bmatrix} \partial_V^2 \mathcal{L} & \partial_V \mathbf{h}^{\mathrm{T}} \\ \partial_V \mathbf{h} & 0 \end{bmatrix} \begin{bmatrix} \delta \mathtt{X}_V \\ \lambda^+ \end{bmatrix} = \begin{bmatrix} -\partial_V f \\ -\mathbf{h} \end{bmatrix}, \tag{8}$$

where $\lambda^+ \triangleq \tilde{\lambda} + \delta\lambda$. The sequence of iterates obtained by taking appropriate steps along $(\delta \mathtt{X}_V, \delta\lambda)$, converges quadratically near a strongly-convex local minimum.

The principal result of this paper is the connection between the unconstrained and constrained formulations,

Theorem 1 (Newton direction) *The Newton direction for unconstrained problem* (4) *is given by an iteration of SQP for the constrained problem* (5)*, when* $\mathtt{X}_{V[\mathcal{G}]}$ *is feasible and when* $\forall v, \lambda_v = \partial_v f$ (12).

Proof. See Appendix A.

Theorem 1 can trivially be extended to arbitrary objective functions on the graph $\mathcal{G}$ and is not restricted to form (5) where objective functions have the same sparsity as the underlying computation. The dual values $\lambda_v = \partial_v f$ can be computed in linear-time with reverse-mode AD.

Theorem 1 immediately leads to the sparse Newton-iteration given in Algorithm 1.

Algorithm 1 Graphical Newton

1: **Input:** initial $\mathtt{X}^0_{\mathtt{Input}}$, tuple $(\mathcal{G}, \{\Phi_v\}, \{\ell_v\})$
2: **repeat**
3: Compute non-inputs, local objectives, and their derivatives (eq. (4)).
4: Set dual-values to $\lambda_v = \partial_v f, \forall v$ by reverse-mode AD (eq. (12)).
5: Solve the KKT-system (eq. (8)).
6: Compute step-length η *via* linesearch on inputs $\mathtt{X}_{\mathtt{Input}}$ (only).
7: Update inputs (only): $\mathtt{X}_{\mathtt{Input}} \leftarrow \mathtt{X}_{\mathtt{Input}} + \eta\delta\mathtt{X}_{\mathtt{Input}}$
8: **until** $||\partial_{\mathtt{Input}} f|| \leq \epsilon$

Algorithm 1 differs markedly from both SQP and current techniques in AD. The non-input primals are computed by running a forward pass on the computation graph, while the dual-values are set to fixed values by running reverse-mode AD. Once the KKT system is solved, the algorithm proceeds to update the primal inputs only, without requiring non-input updates, dual updates, contraint penalties, or any of the other machinery from constrained optimization.

The efficiency of Algorithm 1 stems from the fact that the KKT matrix is typically much sparser than the Hessian. It can be further shown that the common Hessian-accumulation/inversion method is equivalent to a particular, generally a (very) suboptimal, pivot-ordering for solving the sparse KKT system.

3.1. KKT complexity

The run-time of the Algorithm 1 depends crucially on the time taken to solve the sparse KKT system (8) at every iteration. The complexity of solving such sparse systems depends inturn on the support-graph of the underlying matrix.

It has been observed that many problems defined locally on graphs, can be solved in linear-time on trees using dynamic programming. These techniques can be extended to general graphs by grouping vertices in such a way as to *mimic* a tree. The size of the largest cluster in this *tree-decomposition* - termed *tree-width* - governs the dominant term in the time-complexity of the resulting overall algorithm. Computing the tree-decomposition with minimal tree-width is NP-hard, but heuristics for finding elimination orderings are often known to do well in practice.

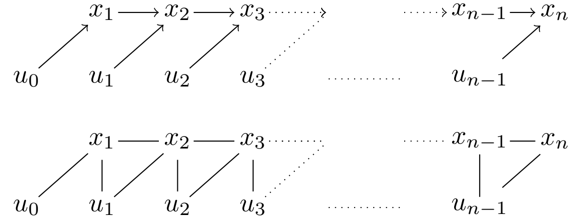

Figure 2. Structure of the optimal control problem as defined in (9). Top: Computational graph. Bottom: Constraint graph.

Run-time complexities in terms of the tree-width are well-established for closely related problems such as Cholesky decomposition, but they appear to be unknown for structured KKT systems such as (8). The following theorem establishes the required bound for structured-KKT systems arising in Algorithm 1.

Theorem 2 (KKT tree-width) *The KKT system* (8) *associated with the constrained problem* (5) *can be solved in time* $O(m\,\mathrm{tw}(\mathcal{G})^3)$*, given the tree-decomposition.*

Proof. See Appendix B.

4. Special cases

4.1. Optimal control

Consider, again, the canonical optimal control problem from (3),

$$\min_{u_0, u_1, \ldots, u_{n-1}} \left[\mathcal{J}(u_0, \ldots, u_{n-1}) \triangleq \sum_{i=0}^{n-1} \ell_i(x_i, u_i) + \ell_n(x_n) \right], \quad \forall i, x_{i+1} \leftarrow \mathtt{f}(x_i, u_i). \tag{9}$$

The computational graph and its moralized relative for this problem are shown in Figure 2. The constraint graph is chordal, and permits multiple optimal elimination orderings.

DYNAMIC PROGRAMMING.

Let's consider the LQR order, $x_n, u_{n-1}, x_{n-1} \ldots, u_0$. The principal minor of the KKT matrix, corresponding to the

Table 1. DDP methods and their key differences.

Method	λ_i	*back-substitution*
DDP (Jacobson & Mayne, 1970)	$\partial_{x_{i+1}} V_{i+1} \cdot \partial_{x_i} \mathtt{f}$	Non-Linear (11)
Stagewise-Newton (De O. Pantoja, 1988)	$\partial_{x_i} \mathcal{J}$	Linear
Nonlinear Stagewise-Newton (Liao & Shoemaker, 1992)	$\partial_{x_i} \mathcal{J}$	Non-Linear (11)
iLQR/iLQG (Li & Todorov)	0	Non-Linear (11)

clique $\{x_n, x_{n-1}, u_{n-1}, \lambda^+_{n-1}\}$ is given by,

$$\left.\begin{array}{cccc} \partial^2_x V_n & & & -\mathtt{I}^\mathtt{T} \\ & \partial^2_x \mathcal{H}_{n-1} & \partial^2_{xu} \mathcal{H}_{n-1} & \partial_x \mathtt{f} \\ & \partial^2_{ux} \mathcal{H}_{n-1} & \partial^2_{uu} \mathcal{H}_{n-1} & \partial_u \mathtt{f} \\ -\mathtt{I}^\mathtt{T} & \partial_x \mathtt{f} & \partial_u \mathtt{f} & 0 \end{array}\right| \begin{array}{c} -\partial_x V_n \\ -\partial_x \ell_{n-1} \\ -\partial_u \ell_{n-1} \\ 0 \end{array},$$

where, $V_n = l_n, \mathcal{H}_{n-1}(x,u) = \ell_{n-1}(x,u) + \lambda_{n-1} \cdot \mathtt{f}(x,u)$. (10)

Note that we can only eliminate variables $\{x_n, u_{n-1}\}$ in the above, since x_{n-1} is part of the separator *i.e* it is connected to nodes outside the clique.

Eliminating $\delta x_n = {\partial^2 V_n}^{-1}(-\partial_x V_n + \lambda^+_{n-1})$,

$$\left.\begin{array}{ccc} \partial^2_x \mathcal{H}_{n-1} & \partial^2_{xu} \mathcal{H}_{n-1} & \partial_x \mathtt{f}^\mathtt{T} \\ \partial^2_{ux} \mathcal{H}_{n-1} & \partial^2_{uu} \mathcal{H}_{n-1} & \partial_u \mathtt{f}^\mathtt{T} \\ \partial_x \mathtt{f} & \partial_u \mathtt{f} & -{\partial^2 V_n}^{-1} \end{array}\right| \begin{array}{c} -\partial_x \ell_{n-1} \\ -\partial_u \ell_{n-1} \\ -{\partial^2 V_n}^{-1} \partial V_n \end{array},$$

Eliminating $\lambda^+_{n-1} = -\partial^2 V_n(-{\partial^2 V_n}^{-1} \partial V_n - \partial_x \mathtt{f} \delta x_{n-1} - \partial_u \mathtt{f}_u \delta u_{n-1})$,

$$\left.\begin{array}{cc} Q_{xx} & Q^\mathtt{T}_{ux} \\ Q_{ux} & Q_{uu} \end{array}\right| \begin{array}{c} -q_x \\ -q_u \end{array},$$

where,

$$Q_{ab} = \partial^2_{ab} \mathcal{H}_{n-1} + \partial_a \mathtt{f}\ \partial^2 V_n\ \partial_b \mathtt{f}^\mathtt{T},$$
$$q_a = \partial_a \ell_{n-1} + \partial_a \mathtt{f}^\mathtt{T}\ \partial V_n.$$

Eliminate $\delta u_{n-1} = {Q_{uu}}^{-1}(-Q_u - Q_{ux} \delta x_{n-1})$,

$$Q_{xx} + Q^\mathtt{T}_{ux} {Q_{uu}}^{-1} Q_{ux} \ \Big|\ -(Q_x + Q^T_{ux} {Q_{uu}}^{-1} Q_u)\ ,$$

Re-write the above as,

$$\partial^2_x V_{n-1} \ \Big|\ -\partial^2_x V_{n-1}.\ ,$$

The eliminations of the adjoining clique can be carried out in a similar manner.

This procedure is identical to the *backward pass* in DDP-methods.[1] In the backsubstitution phase, also known as *forward pass*, the non-linearities can be used directly without having to resort to the use of linearizations,

$$\begin{aligned} \delta x_n &= {\partial^2 V_n}^{-1}(-\partial_x V_n + \lambda^+_{n-1}) \\ &= {\partial^2 V_n}^{-1}(-\partial_x V_n - \partial^2 V_n(-{\partial^2 V_n}^{-1} \partial V_n \\ &\qquad - \partial_x \mathtt{f} \delta x_{n-1} - \partial_u \mathtt{f}_u \delta u_{n-1})) \\ &= \partial_x \mathtt{f} \delta x_{n-1} + \partial_u \mathtt{f}_u \delta u_{n-1}. \\ &\approx \mathtt{f}(x_{n-1} + \delta x_{n-1}, u_{n-1} + \delta u_{n-1}) - \mathtt{f}(x_{n-1}, u_{n-1}). \end{aligned} \quad (11)$$

The Table 1 illustrates how variations in *forward* & *backward* passes, correspond to known DP algorithms. The theory also makes it trivial to generalize these algorithms to higher-order dynamics and constrained problems (Srinivasan & Todorov, 2015).

[1]These computations can be carried out more efficiently & stably in the square-root form.

5. Discussion

The method presented in this paper can be used to compute the Newton step in time $O(m\,\mathrm{tw}^3)$, where tw, m, are the width and size of a tree-decomposition of the computation graph. The method generalizes many specialized DDP algorithms in numerical optimal-control.

However, while the work presented here provides, in a sense, optimal iteration-complexities, many real-world problems in machine-learning also have large tree-widths. The KKT-matrix factorization of such problems is also quite rife with redundant computation, and indicates the necessity for partial symbolic-condensation in sparse LDL (and QP) solvers. These are topics for future work.

References

De O. Pantoja, J. Differential Dynamic Programming and Newton's Method. *International Journal of Control*, 47 (5):1539–1553, 1988.

Jacobson, D. H. and Mayne, D. Q. *Differential Dynamic Programming*. North-Holland, 1970.

Li, W. and Todorov, E. Iterative Linear Quadratic Regulator Design for Nonlinear Biological Movement Systems.

Liao, L.-z. and Shoemaker, C. A. Advantages of Differential Dynamic Programming Over Newton's Method for Discrete-Time Optimal Control Problems. Technical report, Cornell University, 1992.

Srinivasan, A. and Todorov, E. Graphical Newton. *arXiv preprint arXiv:1508.00952*, 2015.

Wright, S. J. Solution of Discrete-Time Optimal Control Problems on Parallel Computers. *Parallel Computing*, 16(2):221–237, 1990.

A. Graphical Newton

This section presents a proof of theorem 1. We proceed by first recalling the $1^{st}, 2^{nd}$-order relations from AD for the objective (4); this is then related to the problem of computing the solutiong to the KKT system in (8) thus completing the proof.

A.1. Derivative relations on $\mathcal{G}$

REVERSE-MODE AD. The first derivative of the objective $f(\cdot)$ (4) can be calculated by applying the chain rule over $\mathcal{G}$,

$$\begin{aligned} \forall v, \quad \partial_v f &= \sum_{s \in v \cup \mathrm{ch}(v)} \partial_v \ell_s + \sum_{d \in \mathrm{ch}(v)} \partial_d f^{\mathsf{T}} \, \partial_v d; \\ v \in \mathrm{pa}(d) &\Rightarrow \partial_v d \triangleq \frac{\partial \Phi_d(\mathtt{X}_{\mathrm{pa}(d)})}{\partial \mathtt{X}_v}. \end{aligned} \tag{12}$$

Since $\mathcal{G}$ is a DAG, there exist child-less nodes (*i.e* $\mathrm{ch}(v) = \emptyset$) from which the recursion can be initialized. The recursion then proceeds backward on $\mathcal{G}$ in a breadth-first order. This algorithm is known in literature as reverse-mode AD.

HESSIAN-VECTOR AD: For a given infinitesimal change $\delta\mathtt{X}_{\mathtt{Input}}$ in the inputs, the first derivatives exhibit a first-order change $\delta[\partial_v f] \triangleq \partial^2_{\mathtt{Input},v} f \cdot \delta\mathtt{X}_{\mathtt{Input}}$, given by the Hessian-vector product. Applying chain-rule over $\mathcal{G}$ again for all terms in (12) we obtain,

$$\begin{aligned} &\forall v, \quad \delta[\partial_v f] = \sum_{s \in v \cup \mathrm{ch}(v)} \left(\sum_{a \in v \cup \mathrm{pa}(s)} \partial^2_{va} \ell_s \cdot \delta\mathtt{X}_a \right) + \\ &\sum_{d \in \mathrm{ch}(v)} \left(\delta[\partial_d f]^{\mathsf{T}} \partial_v d + \sum_{a \in \mathrm{pa}(d)} (\partial_d f^{\mathsf{T}} \, \partial^2_{va} d) \cdot \delta\mathtt{X}_a \right); \\ &\text{where,} \\ &\forall a, \quad \delta\mathtt{X}_a = \sum_{d \in \mathrm{pa}(a)} \partial_d a \cdot \delta\mathtt{X}_d. \end{aligned} \tag{13}$$

These equations can be solved, for a given $\delta\mathtt{X}_{\mathtt{Input}}$ by a forward-backward recursion similar to the one used for solving (12). Computing the Hessian-vector product in this manner takes time $O(|E[\hat{\mathcal{G}}]|\omega(\hat{\mathcal{G}})^2)$, where $\omega(\hat{\mathcal{G}})$ is the clique number of the moralization of $\mathcal{G}$.

A.2. Newton direction

Computing the Newton step requires inverting the Hessian-vector AD process: find $\delta\mathtt{X}_{\mathtt{Input}}$ such that, $\delta[\partial_{\mathtt{Input}} f] = -\partial_{\mathtt{Input}} f$. The inversion of these relations is related to the computation of the SQP direction via Theorem 1,

Theorem 1 (Newton direction) *The Newton direction for unconstrained problem* (4) *is given by an iteration of SQP*

for the constrained problem (5)*, when* $\mathtt{X}_{V[\mathcal{G}]}$ *is feasible and when* $\forall v, \lambda_v = \partial_v f$ (12).

Proof. The second equation in (13) is equivalent to $\partial_V \Phi \cdot \delta S_V = -\Phi$, in (8). Rearranging the first equation from (13), and setting $\delta[\partial_v f] = -\partial_v f$ for all inputs, we obtain $\forall v$,

$$0 = \sum_{\substack{s \in v \cup \mathrm{ch}(v), \\ a \in v \cup \mathrm{pa}(s)}} \partial^2_{va} \ell_s \; \delta \mathtt{X}_a + \sum_{\substack{d \in \mathrm{ch}(v), \\ a \in \mathrm{pa}(d)}} (\partial_d f^{\mathsf{T}} \; \partial^2_{va} d) \; \delta \mathtt{X}_a + \\ - \left(\begin{cases} \delta[\partial_v f] & \mathrm{pa}(v) \neq \emptyset \\ -\partial_v f & \text{otherwise} \end{cases} \right) + \sum_{d \in \mathrm{ch}(v)} (\partial_v d)^{\mathsf{T}} \delta[\partial_d f]. \tag{14}$$

Similarly, expanding the top block in (8) using the definitions in (4) & (8), we obtain $\forall v$,

$$-\partial_v \mathcal{L} = \sum_{\substack{s \in v \cup \mathrm{ch}(v), \\ a \in v \cup \mathrm{pa}(s)}} \partial^2_{va} \ell_s \; \delta \mathtt{X}_a + \sum_{\substack{d \in \mathrm{ch}(v), \\ a \in \mathrm{pa}(d)}} (\lambda_d^{\mathsf{T}} \; \partial^2_{av} d) \; \delta \mathtt{X}_a + \\ - \left(\begin{cases} \delta \lambda_v & \mathrm{pa}(v) \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \right) + \sum_{d \in \mathrm{ch}(v)} (\partial_v d)^{\mathsf{T}} \; \delta \lambda_d, \tag{15}$$

where,

$$\partial_v \mathcal{L} = \begin{cases} \sum_{s \in v \cup \mathrm{ch}(v)} \partial_v \ell_s + \sum_{d \in \mathrm{ch}(v)} \lambda_d \cdot \partial_v d, & \mathrm{pa}(v) = \emptyset \\ \begin{array}{l} \sum_{s \in v \cup \mathrm{ch}(v)} \partial_v \ell_s + \sum_{d \in \mathrm{ch}(v)} \lambda_d \cdot \partial_v d \\ -\lambda_v \end{array}, & \text{otherwise} \end{cases} \tag{16}$$

The result follows from equations (12), (14), (15) & (eq. (16)).

□

B. KKT complexity

In this section, we provide a Message Passing [MP] algorithm for solving KKT systems arising in Algorithm 1 and show that it has a theoretical run-time bound of $O(m \, \mathrm{tw}^3)$, given a tree-decomposition.

B.1. Hypergraph structured QPs

For a hypergraph $\mathcal{H}$, denote the adjacency and incidence matrices by $\mathcal{A}[\mathcal{H}]$ & $\mathcal{B}[\mathcal{H}]$ respectively,

$$\mathcal{A}[\mathcal{H}] \in \mathbb{R}^{|V[\mathcal{H}]| \times |V[\mathcal{H}]|}, \qquad \mathcal{B}[\mathcal{H}] \in \mathbb{R}^{|E[\mathcal{H}]| \times |V[\mathcal{H}]|}, \\ \mathcal{A}[\mathcal{H}]_{uv} = \begin{cases} 1 & \exists e \in E[\mathcal{H}], \; u, v \in e \\ 0 & \text{otherwise} \end{cases} \\ \mathcal{B}[\mathcal{H}]_{eu} = \begin{cases} 1 & u \in e \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

Given such a hypergraph $\mathcal{H}$, the family of QPs we're interested in solving is the following,

$$\min_x \sum_{e \in E[\mathcal{H}]} \frac{1}{2} \mathtt{X}_e^{\mathsf{T}} Q_e \mathtt{X}_e - b_e^{\mathsf{T}} \mathtt{X}_e, \\ \forall e \in E[\mathcal{H}], \quad G_e \mathtt{X}_e = h_e. \tag{18}$$

Assuming that the QP has a bounded solution and that the constraints are full rank, the minimizer to (eq. (18)) is given by the solution to the following KKT system,

$$\begin{bmatrix} Q & G^{\mathsf{T}} \\ G & 0 \end{bmatrix} \begin{bmatrix} x \\ \lambda \end{bmatrix} = \begin{bmatrix} b \\ h \end{bmatrix}, \\ x, b \in \mathbb{R}^{|V|}, \quad \lambda, h \in \mathbb{R}^M, \tag{19}$$

where Q, G, λ, x, b are concatenation of terms defined in (18) respectively. The sparsity/support of (eq. (19)) is closely related to $\mathcal{H}$ because the quadratic part of the KKT equation has the sparsity of the adjacency matrix, and the row of the constraint $G_{i,:}$ has the same sparsity as some edge $e \in E[\mathcal{H}]$.

Tree decomposition: Extending the notion of Dynamic Programming to non-trees (including Hypergraphs) requires a partitioning of the graph so as to satisfy a *lifted* notion of being a tree. Tree decomposition captures the essence of such graph partitions,

Definition 1 *(Tree decomposition) A tree-decomposition of a hypergraph* $\mathcal{H}$ *consists of a tree* $\mathcal{T}$ *and a map* $\chi : V[\mathcal{T}] \to 2^{V[\mathcal{H}]}$*, such that,*

i (Vertex cover) $\cup_{i \in V[\mathcal{T}]} \chi(i) = V[\mathcal{H}]$.

ii (Edge cover) $\forall e \in E[\mathcal{H}], \; \exists i \in V[\mathcal{T}], e \subset \chi(i)$.

iii (Induced sub-tree) $\forall u \in V[\mathcal{H}], \; \mathcal{T}_u \triangleq \mathcal{T}[\{i \in V[\mathcal{T}] | u \in \chi(i)\}]$ *is a non-empty subtree*

The tree-width of a tree-decomposition $\mathcal{T}$ *is defined to be* $\mathrm{tw}(\mathcal{T}) = \max_{v \in V[\mathcal{T}]} |\chi(v)| - 1$. *The tree-width of a graph* $\mathcal{H}$ *is defined to be the minimal tree-width attained by any tree-decomposition of* $\mathcal{H}$.

We define the vertex-induced subgraph in what follows to be $\mathcal{H}[S] \triangleq (V[\mathcal{H}], \{e \cap S, e \in E[\mathcal{H}]\})$. The following lemma ensures that such a decomposition ensures *local dependence*.

Lemma 1 *(Edge separation) Deleting the edge* $xy \in E[\mathcal{T}]$*, renders* $\mathcal{H}[V \backslash (\chi(x) \cap \chi(y))]$ *disconnected.*

Hypertree structured QP: The tree-decomposition itself can be considered a Hypergraph, $(V[\mathcal{H}], \{\chi(u), \forall u \in$

$V[\mathcal{T}]\}$). Such a *Hypertree*[2] can also be thought of as a chordal graph. We assume henceforth that the given graph $\mathcal{H}$ is a hypertree, and that $\mathcal{T}$ is its tree-decomposition.

A message passing [MP] algorithm for solving (19) on such Hypertrees is given below. The gather stage of the message passing algorithm is illustrated in algorithm 2 [3]. The function, Factorize, computes the partial LU decomposition of its arguments; we describe below, its operation.

Denote the vertices that are interior to l by $\iota = \chi(l) \cap \chi(p)$, and those on the boundary (*i.e* common to p, l) by $\xi = \chi(l) \backslash \chi(p)$, and let $r = \operatorname{rank}(\tilde{Q}_{\iota,\iota})$. The function computes Gaussian-BP messages from block pivots $\{2, 3\}$ to $\{1, 4\}$ in (20). Note that, unlike Gaussian-BP, the matrices in (20) are not necessarily positive definite, but are however invertible.

Algorithm 2 Graphical QP

1: **Given**: $\mathcal{T}, \mathcal{H}, \{Q_e\}, \{b_e\}, \{G_e\}, \{h_e\}$.
2:
3: **function** GatherMessage($l, p, \mathcal{T}$)
4: $(\tilde{Q}_l, \tilde{b}_l, \tilde{G}_l, \tilde{h}_l) \leftarrow (Q_l, b_l, G_l, h_l)$
5: **for** $c \in \delta_{\mathcal{T}}(l) \backslash p$ **do**
6: $\quad (Q_{c\to l}, G_{c\to l}, b_{c\to l}, h_{c\to l}) \leftarrow$ GatherMessage($c, p, \mathcal{T}$)
7: $\quad (\tilde{Q}_l, \tilde{b}_l) \leftarrow (\tilde{Q}_l, \tilde{b}_l) + (Q_{c\to l}, b_{c\to l})$
8: $\quad \tilde{G}_l \leftarrow [\tilde{G}_l; G_{c\to l}], \tilde{h}_l \leftarrow [\tilde{h}_l; h_{c\to l}]$
9: **end for**
10: **return** Factorize($\chi(l), \chi(p), \tilde{Q}_l, \tilde{b}_l, \tilde{G}_l, \tilde{h}_l$)
11:
12: **function** Factorize($\chi(l), \chi(p), \tilde{Q}, \tilde{b}, \tilde{G}, \tilde{h}$)
13: $(\xi, \iota) \leftarrow (\chi(l) \backslash \chi(p), \chi(l) \cap \chi(p))$
14: $r \leftarrow \operatorname{rank}(\tilde{Q}_{\iota,\iota})$
15: **return** Gaussian-BP messages from (20).
16: **return**

$$\begin{bmatrix} \tilde{Q}_{\xi\xi} & \tilde{Q}_{\iota\xi}^{\mathrm{T}} & \tilde{G}_{:r,\xi}^{\mathrm{T}} & \tilde{G}_{r:,\iota}^{\mathrm{T}} \\ \tilde{Q}_{\iota\xi} & \tilde{Q}_{\iota\iota} & \tilde{G}_{:r,\iota}^{\mathrm{T}} & \tilde{G}_{r:,\iota}^{\mathrm{T}} \\ \tilde{G}_{:r,\xi} & \tilde{G}_{:r,\iota} & 0 & 0 \\ \tilde{G}_{r:,\xi} & \tilde{G}_{r:,\iota} & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathtt{x}_\xi \\ \mathtt{x}_\iota \\ \lambda_{:r} \\ \lambda_{r:} \end{bmatrix} = \begin{bmatrix} \tilde{b}_\xi \\ \tilde{b}_\iota \\ \tilde{h}_{:r} \\ \tilde{h}_{r:} \end{bmatrix} \tag{20}$$

Gaussian Belief-Propagation is essentially a re-statement of LU decomposition, and consists of messages of the form,

$$\mu_{i\to j} := [J_{i\to j}, h_{i\to j}] = [J_{ii}, h_i] - \sum_{k\in\delta(i)\backslash j} J_{ik} J_{k\to i}^{-1} [J_{ki}, h_{k\to i}],$$
$$\mu_i = J_{i\to j}^{-1}(h_{i\to j} - J_{ij}\mu_j), \tag{21}$$

where $J\mu = h$ is the equation that is to be solved. These can be replaced by appropriate square-root forms to obtain instead, an LDL decomposition.

Theorem 3 *The linear equation* (19) *can be solved in time* $O(|\mathcal{H}| \operatorname{tw}(\mathcal{H})^3)$*, given the tree-decomposition,* via *Algorithm 2.*

Proof. The correctness of the algorithm follows from Lemma 1. The bound holds trivially if, $\operatorname{rank} \tilde{G} \leq \operatorname{rank} \tilde{Q}_{\iota,\iota}$, at every step of the algorithm. Otherwise, by realizing that $\tilde{G}_{l\to p}$, can't have rank more than $|\chi(p)|$, the proof follows. □

[2] There are multiple definitions of a *Hypertree*; we use the term to mean a maximal Hypergraph, whose tree-decomposition can be expressed in terms of its edges.

[3] Note that the addition is performed vertex label-wise in Line 6 of Algorithm algorithm 2.