

Parallel ADMM for robust quadratic optimal resource allocation problems

Zawar Qureshi¹

Sebastian East¹

Mark Cannon¹

Abstract—An alternating direction method of multipliers (ADMM) solver is described for optimal resource allocation problems with separable convex quadratic costs and constraints and linear coupling constraints. We describe a parallel implementation of the solver on a graphics processing unit (GPU) using a bespoke quartic function minimizer. An application to robust optimal energy management in hybrid electric vehicles is described, and the results of numerical simulations comparing the computation times of the parallel GPU implementation with those of an equivalent serial implementation are presented.

I. INTRODUCTION

An important class of control problems can be formulated as optimal resource allocation problems with separable convex cost functions, separable convex constraints and linear coupling constraints. Examples include energy management in hybrid electric vehicles [1], [2], optimal power flow and economic dispatch problems [3], portfolio selection [4], and diverse problems in manufacturing, production planning and finance [5]. The computation required to solve these problems grows with the length of the horizon over which performance is optimized, and the rate of growth becomes more pronounced for scenario programming formulations [6] that provide robustness to demand uncertainty. This limits the applicability of strategies such as model predictive control that require online solutions of resource allocation problems.

The alternating direction method of multipliers (ADMM) is a first order optimization method that can exploit separable cost and constraint functions, thus forming the basis of a solver suitable for parallel implementation [7], [8]. This paper describes a graphics processing unit (GPU) parallel implementation of ADMM that avoids the growth in computation time with horizon that is observed in serial implementations. This is possible because GPUs have large numbers (typically thousands) of processing units and efficient memory architectures. Parallel ADMM implemented on general purpose GPUs therefore provides a low-cost approach for time-critical optimal resource allocation problems.

Under the assumption that the separable cost and constraint functions are defined by quadratic polynomials, we derive in this paper an ADMM iteration in which the most computationally demanding step is the minimization of a set of quartic polynomials. Iterative methods have previously been proposed for solving polynomial functions on the GPU [9], [10]. However, to retain the convergence guarantees of ADMM, we devise an exact quartic minimization

algorithm for GPUs. The solver was implemented on the GPU using CUDA [11] and we compare its computation times with an equivalent serial implementation for a range of problem sizes. We also describe the application of ADMM to robust optimal energy management in a plugin hybrid electric vehicle (PHEV). The GPU implementation of this solver was tested in simulations against corresponding serial implementations. The results demonstrate a significant speed-up and indicate that the approach is feasible for robust supervisory control of PHEVs with realistic horizon lengths.

II. PROBLEM STATEMENT

Consider a resource allocation problem in which a collection of resources, which could for example represent alternative power sources onboard a hybrid vehicle, is to be used to meet predicted demand at each discrete time instant over a future horizon. For a horizon of n time steps, let y_k denote the demand and $x_k^{(i)}$ denote the resource allocated from source i at time step k , for $k \in \mathbb{N}_n := \{1, \dots, n\}$ and $i \in \mathbb{N}_m := \{1, \dots, m\}$. In order to meet the predicted demand, we require

$$\sum_{i=1}^m x_k^{(i)} \geq y_k, \quad k \in \mathbb{N}_n$$

where $x_k^{(i)}$ is subject to bounds, $x_k^{(i)} \in [\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]$, $k \in \mathbb{N}_n$ as well as constraints on total capacity:

$$\sum_{k=1}^n g_k^{(i)}(x_k^{(i)}) \leq c^{(i)}, \quad i \in \mathbb{N}_m.$$

Capacity constraints represent, for example, the total energy, $c^{(i)}$, available from the i th power source over the n -step horizon, with losses accounted for through the functions $g_k^{(i)}$.

The cost associated with $x_k^{(i)}$ is denoted $f_k^{(i)}(x_k^{(i)})$. This represents, for example, the operating cost or fuel cost (including losses) associated with the i th resource at time k . The objective is to minimize total cost subject to constraints:

$$\begin{aligned} \min_{x_k^{(i)} \in [\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]} & \sum_{i=1}^m \sum_{k=1}^n f_k^{(i)}(x_k^{(i)}) \\ \text{subject to} & \sum_{i=1}^m x_k^{(i)} \geq y_k \quad k \in \mathbb{N}_n \\ & \sum_{k=1}^n g_k^{(i)}(x_k^{(i)}) \leq c^{(i)} \quad i \in \mathbb{N}_m. \end{aligned} \quad (1)$$

This problem is convex if $f_k^{(i)}$ and $g_k^{(i)}$ are convex functions of their arguments. This is a common assumption in energy management problems with nonlinear losses [12], [13].

¹Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK.

qureshizawar@gmail.com, sebastian.east@eng.ox.ac.uk, mark.cannon@eng.ox.ac.uk

Assumption 1: The functions $f_k^{(i)}$ and $g_k^{(i)}$ are convex, for all $i \in \mathbb{N}_m$ and $k \in \mathbb{N}_n$.

If the demand sequence $y = \{y_k\}_{k \in \mathbb{N}_n}$ is stochastic, we assume that independent samples $y^{(j)}$, $j \in \mathbb{N}_q := \{1, \dots, q\}$, can be drawn from the probability distribution of y .

Assumption 2: The samples $y^{(j)} = \{y_k^{(j)}\}_{k \in \mathbb{N}_n, j \in \mathbb{N}_q}$ are such that $y^{(j)}$ and $y^{(l)}$ are independent for $l \neq j$.

Various robust formulations of the nominal allocation problem (1) are possible using a set of sampled sequences $y^{(j)}$, $j \in \mathbb{N}_q$. For example, an open loop formulation optimizes a single sequence $\{x_k^{(i)}\}_{k \in \mathbb{N}_n}$ for each $i \in \mathbb{N}_m$ with the predicted power demand sequence in (1) replaced by $\{\max_j y_k^{(j)}\}_{k \in \mathbb{N}_n}$. Alternatively, a receding (or shrink-ing) horizon approach introduces feedback by solving the resource allocation problem at each sampling instant using current capacity constraints and demand estimates. In this case a sequence $\{x_k^{(i,j)}\}_{k \in \mathbb{N}_n}$ is computed for each $i \in \mathbb{N}_m$ and $j \in \mathbb{N}_q$ subject to the additional constraint that, for each i , the first element in the sequence should be equal for all j :

$$\begin{aligned} \min_{x_1^{(i)} \in \mathbb{R}, x_k^{(i,j)} \in [\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]} \quad & \frac{1}{q} \sum_{j=1}^q \sum_{i=1}^m \sum_{k=1}^n f_k^{(i,j)}(x_k^{(i,j)}) \\ \text{subject to} \quad & \sum_{i=1}^m x_k^{(i,j)} \geq y_k^{(j)} \quad k \in \mathbb{N}_n, j \in \mathbb{N}_q \\ & \sum_{k=1}^n g_k^{(i,j)}(x_k^{(i,j)}) \leq c^{(i)} \quad i \in \mathbb{N}_m, j \in \mathbb{N}_q \\ & x_1^{(i,j)} = x_1^{(i)} \quad i \in \mathbb{N}_m, j \in \mathbb{N}_q \end{aligned} \quad (2)$$

Although the objective in (2) is an empirical mean, the approach of this paper allows for alternative definitions such as the maximum over $j \in \mathbb{N}_q$ of $\sum_{i,k} f_k^{(i,j)}(x_k^{(i,j)})$. The probability that the solution of (2) violates the constraints of (1) has a known probability distribution that depends on the number, q of samples and on the structural properties of (2) – see e.g. [6], [14] for details – moreover these results apply for any demand probability distribution.

An ADMM solver for (2) is summarised in Appendix A. The approach exploits the separable nature of the cost and constraints of (2), which are sums of functions that depend only on variables at individual time-steps, to derive an iteration in which the greatest contribution to the overall computational load is the minimization of scalar polynomials in (7a). Consequently the iteration does not require matrix operations (inversions or factorizations) and, as discussed below, it is suitable for implementation on extremely simple computational hardware. For each $i \in \mathbb{N}_m$, the polynomial minimization in step (7a) can be performed in parallel for $k \in \mathbb{N}_n$ and $j \in \mathbb{N}_q$, since the update for $x_k^{(i,j)}$ does not depend on $x_l^{(i,r)}$ for $(l, r) \neq (k, j)$. A parallel implementation therefore allows the ADMM iteration to be computed quickly even for long horizons and large numbers of scenarios.

In the following sections, we describe a GPU-based parallel implementation of the ADMM iteration defined in (7a-i). We make the assumption that $f_k^{(i,j)}$ and $g_k^{(i,j)}$ are convex quadratic functions of their arguments. Under

this assumption (2) becomes a quadratically constrained quadratic program, and the iteration (7a) then requires the minimization of a 4th order polynomial, for each i, j and k .

Assumption 3: $f_k^{(i,j)}$ and $g_k^{(i,j)}$ are quadratic functions:

$$\begin{aligned} f_k^{(i,j)}(x) &= \alpha_{k,2}^{(i,j)} x^2 + \alpha_{k,1}^{(i,j)} x + \alpha_{k,0}^{(i,j)} \\ g_k^{(i,j)}(x) &= \beta_{k,2}^{(i,j)} x^2 + \beta_{k,1}^{(i,j)} x + \beta_{k,0}^{(i,j)} \end{aligned}$$

for $i \in \mathbb{N}_m$, $j \in \mathbb{N}_q$, $k \in \mathbb{N}_n$.

III. PARALLELIZATION

This section discusses parallelization of the ADMM iteration in Appendix A, focussing on the minimization step (7a). Under Assumption 3, this requires the minimization of 4th order polynomials, for which we describe a tailored algorithm in Section III-B. We consider NVIDIA GPUs programmed in CUDA C/C++ [11] based on a single instruction multiple thread (SIMT) execution model. The CPU (host) and GPU (device) have separate memory, making it necessary to copy data between the host and device memory. Copying data back and forth adversely affects computation time, and the computational benefits of parallelization may therefore only become apparent for large problems. We therefore investigate in Section III-C the dependence of the computational speed-up that can be achieved via parallelization on the number of quartic polynomials to be minimized.

A. CUDA programming model

An algorithm involving N parallel executions of a function (known as a *kernel*) is implemented by executing N separate CUDA threads. Each CUDA thread is assigned an index within a one-, two- or three-dimensional block of threads. This provides a means of invoking computation across the elements of a vector, matrix, or volume. On current NVIDIA GPUs, a thread block may contain up to 1024 threads. This limit is due to the requirement that all threads of a block must reside on the same processor core and share the limited memory resources of that core. A kernel can be executed by multiple equally-shaped thread blocks, which are themselves organized into a grid of thread blocks.

During execution, CUDA threads may access data from multiple memory spaces. Each thread has private local memory; each thread block has shared memory visible to all threads of the block and with the same lifetime as the block; and all threads have access to the same global memory.

For the CUDA implementations discussed in Sections III-C and IV-B, individual kernels were devised for specific elementwise operations while cuBLAS library functions were used for vector-matrix multiplications. To compare the performance of non-parallel implementations, serial algorithms were written in C++ with compiler optimizations (e.g. vectorization where possible) set for speed (/Ox) unless stated otherwise. Intel Math Kernel Library BLAS routines were used for vector-matrix multiplications. Both implementations were coded using Microsoft Visual Studio.

B. Quartic minimization

Iterative solution methods, such as Newton's method [9] and recursive de Casteljau subdivision [10] have previously been proposed for solving systems of polynomials on a GPU. However, when used to minimize the quartic polynomial in (7a), iterative methods can adversely affect ADMM convergence due to residual errors in their solutions. For parallel GPU implementations of the iteration (7) therefore, a better approach is to compute the closed form solutions of the cubic polynomial obtained after differentiating the quartic in (7a) with respect to $x_k^{(i,j)}$. Algebraic solutions of cubic equations can be derived in various ways, including Cardano's method and Vieta's method e.g. [15, Sec. 14.7]). We use a combination of these methods to first determine whether roots are real and/or repeated before they are calculated. For problems in which real solutions are expected, this allows significant computational savings and avoids the memory requirements of storing complex roots.

Consider the quartic polynomial

$$f(x) = Ax^4 + Bx^3 + Cx^2 + Dx + E \quad (3)$$

and let $b = 3B/4A$, $c = C/2A$ and $d = D/4A$. The extremum points of f are the solutions of the cubic equation

$$x^3 + bx^2 + cx + d = 0. \quad (4)$$

Let Q , R and Δ be defined by

$$Q = c/3 - b^2/9, \quad R = bc/6 - b^3/27 - d/2 \\ \Delta = Q^3 + R^2.$$

If $\Delta \leq 0$, all three roots of (4) are real. In this case let

$$\theta = \cos^{-1}\left(R/\sqrt{-Q^3}\right),$$

then the roots are

$$x_a = 2\sqrt{-Q} \cos\left(\theta/3\right) - b/3 \\ x_b = 2\sqrt{-Q} \cos\left(\theta/3 + 2\pi/3\right) - b/3 \\ x_c = 2\sqrt{-Q} \cos\left(\theta/3 + 4\pi/3\right) - b/3.$$

If $\Delta > 0$, there is only one real root. In this case, defining

$$S = \sqrt[3]{R + \sqrt{\Delta}}, \quad T = \sqrt[3]{R - \sqrt{\Delta}},$$

the real root is given by

$$x_a = S + T - b/3.$$

For the special case of $Q = R = 0$, the three roots are real and repeated, and Vieta's formula gives the solution

$$x_a = x_b = x_c = -b/3.$$

We additionally require the solver to determine which root of (4) is the minimizing argument of f . This is simplified by sorting the roots and discarding the middle root, which is necessarily a maximizer. The minimizer can then be determined by comparing the value of the objective function at the two remaining roots. The quartic minimization algorithm is summarised in Algorithm 1.

Algorithm 1: Quartic minimizer

Input : Coefficients A, B, C, D

Output: Minimizer x^*

```

1 Evaluate  $b, c, d, Q, R, \Delta$ ;
2 if  $\Delta > 0$  then
3   Evaluate  $S, T$ ;
4    $x^* \leftarrow S + T - b/3$ ;
5 else if  $Q = R = 0$  then
6    $x^* \leftarrow -b/3$ ;
7 else
8   Evaluate  $x_a, x_b, x_c$ ;
9    $(x_1, x_2, x_3) \leftarrow \text{sort}(x_a, x_b, x_c)$ ;
10   $\delta f \leftarrow \frac{1}{4}(x_1^4 - x_3^4) + \frac{b}{3}(x_1^3 - x_3^3) + \frac{c}{2}(x_1^2 - x_3^2) + d(x_1 - x_3)$ ;
11  if  $\delta f > 0$  then
12     $x^* \leftarrow x_3$ ;
13  else
14     $x^* \leftarrow x_1$ ;

```

C. Comparison of CPU/GPU performance

The results presented in this section and Section IV-B were obtained using an Intel i5-7300HQ CPU with base clock speed 2.5 GHZ, 12GB of RAM, and an Nvidia GTX 1060 3GB GPU. All operations were performed with single-precision floating point numbers.

To test performance we generate N quartic equations with randomly chosen coefficients. For the CPU implementation we solve the N equations sequentially using Algorithm 1. For the GPU implementation, all available threads (up to a limit of 18432 threads for the Nvidia GTX 1060 3GB) are scheduled to perform Algorithm 1 in parallel. The test procedure for the CPU-based algorithm was as follows.

- 1) Generate N random sets of quartic coefficients.
- 2) Record start time T_1 .
- 3) Minimize the N quartics using Algorithm 1.
- 4) Record end time T_2 and elapsed time $T = T_2 - T_1$.

The test procedure for the GPU parallel cubic solver is similar to the CPU procedure with step 3 changed to:

- 3) a) Input data to the GPU.
- b) Minimize the N quartics using Algorithm 1 implemented on the GPU.
- c) Copy results back to the CPU.

The test procedure was repeated at least 10 times and the average value taken for the run-time. Tests were carried out with and without CPU compiler speed optimizations.

Figure 1 compares the GPU and CPU execution speeds as N varies. For $N > 10^6$, a 30-fold increase in speed is obtained using the GPU in comparison with the CPU implementation with 4 cores. The speed-up factor rises to 210 compared with a single-core CPU implementation with speed optimizations off. Although the CPU implementation is faster for $N < 10^3$, the GPU becomes significantly faster for large N . This is expected as the number of GPU threads must be sufficiently large to counter the effects of overheads such as memory transfer times.




Fig. 1. Speed comparison for minimizing N quartic equations. For CPU implementations, speed optimizations were either enabled or disabled using compiler flags `/Ox` and `/Od` respectively.

IV. PHEV SUPERVISORY CONTROL

This section describes an application of the robust resource allocation problem (2) and the parallelization approach of Section III to optimal energy management for a Plug-in Hybrid Electric Vehicle (PHEV). PHEVs employ hybrid powertrains that allow the use of more than one energy source. Here we consider a PHEV with two power sources: an electric motor/generator with battery storage, and an internal combustion engine. The supervisory control problem involves determining the optimal power split between the internal combustion engine and the electric motor in order to minimize a performance cost, defined here as the fuel consumption over a prediction horizon, while ensuring that the current and predicted future power demand is met and that battery stored energy remains within allowable limits.

We consider a parallel hybrid architecture [1], [2] in which the propulsive power (y) demanded from the vehicle's powertrain is provided by the sum of power from the combustion engine ($x^{(1)}$) and power from the electric motor ($x^{(2)}$). Independent samples, $\{y_k^{(j)}\}_{j \in \mathbb{N}_q, k \in \mathbb{N}_n}$, of the future demand sequence are assumed to be available, and the power balance

$$y_k^{(j)} \leq x_k^{(1,j)} + x_k^{(2,j)}, \quad j \in \mathbb{N}_q, \quad k \in \mathbb{N}_n$$

is imposed over an n -step future horizon. For each demand scenario j and time step k the engine speed $\omega_k^{(j)}$ is assumed to be known and equal to the electric motor speed whenever the clutch is engaged.

The fuel energy consumed during the k th time interval is obtained from a quasi-static map, $f_k^{(1,j)}(x_k^{(1,j)})$, relating engine power output to fuel consumption for the given engine speed $\omega_k^{(j)}$. Likewise, the battery energy output during the k th period, $g_k^{(2,j)}(x_k^{(2,j)})$, is determined for a given $\omega_k^{(j)}$ by a quasi-static model of battery and electric motor losses. The energy stored in the battery at the end of the horizon must be no less than a given terminal quantity, E_n , and hence the fall in battery stored energy over the prediction horizon must

not exceed $\Delta E = E_0 - E_n$, where E_0 is the initial battery state. No constraint is applied to total fuel usage and no cost associated with electrical energy usage (i.e. $g_k^{(1,j)} = 0$ and $f_k^{(2,j)} = 0$), so the optimization problem (2) becomes

$$\begin{aligned} \min_{x_1^{(i)} \in \mathbb{R}, x_k^{(i,j)} \in [\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]} \quad & \frac{1}{q} \sum_{j=1}^q \sum_{k=1}^n f_k^{(1,j)}(x_k^{(1,j)}) \\ \text{subject to} \quad & x_k^{(1,j)} + x_k^{(2,j)} \geq y_k^{(j)}, \quad j \in \mathbb{N}_q, \quad k \in \mathbb{N}_n \\ & \sum_{k=1}^n g_k^{(2,j)}(x_k^{(2,j)}) \leq \Delta E, \quad j \in \mathbb{N}_q \\ & x_1^{(i,j)} = x_1^{(i)}, \quad i = 1, 2, \quad j \in \mathbb{N}_q \end{aligned} \quad (5)$$

The updates of $x_k^{(i,j)}$ in (7a) therefore require the minimization of a quadratic for $i = 1$ and a quartic polynomial for $i = 2$. The bounds $\underline{x}_k^{(i)}$ are determined by constraints on power and torque output of powertrain components (see e.g. [2]); for example engine braking ($x_k^{(1,j)} < 0$) is ignored, so $\underline{x}_k^{(1)} = 0$. Assumption 1 (convexity) is satisfied if $f_k^{(1,j)}$ and $g_k^{(2,j)}$ are convex functions, which is usually the case for real powertrain loss maps ([2]). In addition, cost and constraint functions are approximated by fitting quadratic functions to $f_k^{(1,j)}$ and $g_k^{(2,j)}$ at each time-step (see [2] for details), and Assumption 3 is therefore satisfied.

At each sampling instant a supervisory control algorithm based on the online solution of (5) with a shrinking horizon can be summarised as follows.

Algorithm 2 (Supervisory PHEV control):

- 1) Draw samples, $\{y_k^{(j)}\}_{j \in \mathbb{N}_q, k \in \mathbb{N}_n}$, from the distribution of the predicted power demand sequence.
- 2) For the current battery energy level, E_0 , set $\Delta E = E_0 - E_n$, and compute for each $j \in \mathbb{N}_q$ the engine speed sequence $\{\omega_k^{(j)}\}_{k \in \mathbb{N}_n}$ corresponding to $\{y_k^{(j)}\}_{k \in \mathbb{N}_n}$, and hence determine $f_k^{(1,j)}$ and $g_k^{(2,j)}$ for $k \in \mathbb{N}_n$.
- 3) Solve (5) by applying the ADMM iteration (7a-i) until the primal and dual residuals satisfy $r < \bar{r}$ and $\sigma < \bar{\sigma}$.
- 4) Apply the engine and electric motor power $x_1^{(1)}, x_1^{(2)}$.

A. Modelling assumptions and optimization parameters

We consider a 1900 kg vehicle with a hybrid powertrain consisting of a 100 kW gasoline internal combustion engine, 50 kW electric motor, and a 21.5 Ah lithium-ion battery. Power demand sequences were generated by adding random sequences (Gaussian white noise of standard deviation 10 kW, low-pass filtered with cutoff frequency 0.02 Hz) to the power demand computed for the FTP-75 drive cycle. The sampling interval was fixed at 1 s. The times at which the vehicle stopped were also modified to simulate variable traffic conditions. This method of generating power demand is used simply to illustrate the computational aspects of the approach; more accurate methods of generating scenarios in this context (e.g. [16], [17]) have been proposed in the literature on stochastic model predictive control.

The initial battery energy level E_0 was set at 60% of battery capacity, E_{\max} , and the terminal level, E_n , at 50% of




Fig. 2. Predicted battery state of charge profiles corresponding to the optimal solution of (5) with 100 power demand scenarios generated from random perturbations of the FTP-75 drive cycle.

capacity. Whenever the power demand $y_k^{(j)}$ is negative, 40% is assumed to be recovered through regenerative braking.

Thresholds for the primal and dual residuals were set at $\bar{r} = 10^{-6} \Delta E$ and $\bar{\sigma} = 10^{-2}$. Step-length parameters $\rho_1, \rho_2, \rho_3, \rho_4$ were adapted at each ADMM iteration according to $\rho_i \leftarrow \tau \rho_i$ if $r/\sigma > 1.2\bar{r}/\bar{\sigma}$ and $\rho_i \leftarrow \rho_i/\tau$ if $r/\sigma < 0.8\bar{r}/\bar{\sigma}$ with $\tau = 1.1$ and initial values $\rho_1 = 10^{-4}$, $\rho_2 = 2 \times 10^{-6}$, $\rho_3 = \rho_4 = 5 \times 10^{-6}$. This adaptive approach, which is based on the discussion in Sec. 3.4.1 of [8], was found to reduce the effect of the value of q on the rate of convergence of the ADMM iteration. A typical solution for the battery state of charge for $q = 100$ demand scenarios is shown in Fig. 2.

B. Comparison of CPU/GPU performance

The CPU-based ADMM implementation was tested using the following procedure:

- 1) Perform steps 1 and 2 of Algorithm 2.
- 2) Record start time T_1 .
- 3) Execute step 3 of Algorithm 2 on the CPU.
- 4) Record end time T_2 and elapsed time $T = T_2 - T_1$.

For the parallel implementation, step 3 was replaced by:

- 3) a) Input data to the GPU.
- b) Execute step 3 of Algorithm 2 on the GPU.
- c) Copy results back to the CPU.

For numbers of demand scenarios q between 5 and 500, Figure 3 shows that GPU runtimes were between 10 and 20 times faster than tests executed on a 4-core CPU with (/Ox) speed optimizations. The number of ADMM iterations required tends to increase with q , although for the chosen parameters a significant reduction is observed around $q = 20$. The speed-up achieved by the GPU implementation increases with q , even though for $q > 10$ the number of threads required to execute the ADMM iteration exceeds the limit that can be executed simultaneously on the GPU. To reduce the effect on overall performance of computing primal and dual residuals, the termination conditions were checked only once every 10 iterations in both implementations.




Fig. 3. Mean runtimes of CPU and GPU (parallel) implementations (left vertical axis) and number of ADMM iterations (right vertical axis) against number of demand scenarios (horizontal axis).

V. CONCLUSIONS

This paper describes an ADMM algorithm for robust quadratic resource allocation problems. We discuss a parallel implementation of the algorithm and its application to optimal energy management for PHEVs with uncertain power demand. Our results suggest that it is feasible to solve optimal PHEV energy management problems online using inexpensive computing hardware. We also note that the state of the art in low-cost parallel processing hardware is evolving rapidly and that next generation GPUs can provide significant speed improvements over the results reported here.

APPENDIX A: ADMM ITERATION

To express (2) in a form suitable for ADMM, we introduce slack variables $s_k^{(j)}$ into the demand constraints and introduce dummy variables $z_k^{(i,j)}, h^{(i,j)}$ into the capacity constraints:

$$\begin{aligned}
 & \min_{\substack{x_1^{(i)} \in \mathbb{R}, x_k^{(i,j)} \in [\underline{x}_k^{(i)}, \bar{x}_k^{(i)}], z_k^{(i,j)} \in \mathbb{R} \\ s_k^{(j)} \geq 0, h^{(i,j)} \leq c^{(i)}}} \frac{1}{q} \sum_{i,j,k} f_k^{(i,j)}(x_k^{(i,j)}) \\
 & \text{subject to } \sum_i x_k^{(i,j)} - y_k^{(j)} = s_k^{(j)} \quad j \in \mathbb{N}_q, k \in \mathbb{N}_n \\
 & \quad g_k^{(i,j)}(x_k^{(i,j)}) = z_k^{(i,j)} \quad i \in \mathbb{N}_n, j \in \mathbb{N}_q, k \in \mathbb{N}_n \\
 & \quad \sum_k z_k^{(i,j)} = h^{(i,j)} \quad i \in \mathbb{N}_n, j \in \mathbb{N}_q \\
 & \quad x_1^{(i,j)} = x_1^{(i)} \quad i \in \mathbb{N}_n, j \in \mathbb{N}_q
 \end{aligned} \tag{6}$$

This formulation preserves separability of the nonlinear cost and constraints by imposing linear constraints on $z_k^{(i,j)}$. Let $\mathcal{I}_{\mathbb{I}}(\cdot)$ denote the indicator function of the interval \mathbb{I} (for scalar x , $\mathcal{I}_{\mathbb{I}}(x) = 0$ if $x \in \mathbb{I}$ and $\mathcal{I}_{\mathbb{I}}(x) = \infty$ if $x \notin \mathbb{I}$; for vector x , $\mathcal{I}_{\mathbb{I}}(x)$ is the vector with k th element equal to $\mathcal{I}_{\mathbb{I}}(x_k)$). We define the augmented Lagrangian function

$$\begin{aligned}
 L(x_1, x, z, s, h, \mu, \lambda, \nu, p) &:= \frac{1}{q} \sum_{i,j,k} f_k^{(i,j)}(x_k^{(i,j)}) \\
 &+ \sum_{i,j} \left[\mathcal{I}_{[\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]}(x_k^{(i,j)}) + \frac{\rho_1}{2} \|z_k^{(i,j)} - g^{(i,j)}(x_k^{(i,j)}) + \lambda^{(i,j)}\|^2 \right]
 \end{aligned}$$

$$\begin{aligned}
& + \sum_{i,j} \left[\mathcal{I}_{(-\infty, c(i)]}(h^{(i,j)}) + \frac{\rho_2}{2} (h^{(i,j)} - \mathbf{1}^\top z^{(i,j)} + p^{(i,j)})^2 \right] \\
& + \sum_j \left[\mathbf{1}^\top \mathcal{I}_{[0, \infty)}(s^{(j)}) + \frac{\rho_3}{2} \|s^{(j)} - \sum_i x^{(i,j)} + y^{(j)} + \mu^{(j)}\|^2 \right] \\
& + \sum_{i,j} \frac{\rho_4}{2} (x_1^{(i)} - x_1^{(i,j)} + \nu^{(i,j)})^2
\end{aligned}$$

Here $\mu^{(j)}$, $\lambda^{(i,j)}$, $\nu^{(i,j)}$, $p^{(i,j)}$ and $\nu^{(i,j)}$ are multipliers for the equality constraints in (6); $y^{(j)}$, $x^{(i,j)}$, $z^{(i,j)}$, $s^{(j)}$ and $g^{(i,j)}(x^{(i,j)})$ are the vectors with k th elements equal to $y_k^{(j)}$, $x_k^{(i,j)}$, $z_k^{(i,j)}$, $s_k^{(j)}$ and $g_k^{(i,j)}(x_k^{(i,j)})$, and $\mathbf{1} = [1 \dots 1]^\top$. Also ρ_1, \dots, ρ_4 are non-negative scalars that control the step length of multiplier updates in the ADMM iteration [8, Sec. 3.1]. These design parameters control the rate of convergence and can be tuned to the problem at hand (see e.g. [18]).

At each iteration, L is minimized with respect to primal variables $x^{(i,j)}$, $x_1^{(i)}$, $z^{(i,j)}$, $s^{(j)}$, $h^{(i,j)}$. The primal and dual variable updates are given, for $k \in \mathbb{N}_n$, $i \in \mathbb{N}_m$, $j \in \mathbb{N}_q$, by

$$\begin{aligned}
x_k^{(i,j)} \leftarrow \Pi_{[\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]} \left(\underset{x_k^{(i,j)}}{\operatorname{argmin}} \left\{ \frac{1}{q} f_k^{(i,j)}(x_k^{(i,j)}) \right. \right. \\
+ \frac{\rho_1}{2} (z_k^{(i,j)} - g_k^{(i,j)}(x_k^{(i,j)}) + \lambda_k^{(i,j)})^2 \\
+ \frac{\rho_3}{2} (s_k^{(j)} - \sum_i x_k^{(i,j)} + y_k^{(j)} + \mu_k^{(j)})^2 \\
\left. \left. + \delta_{k,1} \frac{\rho_4}{2} (x_1^{(i)} - x_1^{(i,j)} + \nu^{(i,j)})^2 \right\} \right) \quad (7a)
\end{aligned}$$

$$\begin{aligned}
z^{(i,j)} \leftarrow g^{(i,j)}(x^{(i,j)}) - \lambda^{(i,j)} \\
+ \frac{\rho_2}{\rho_1 + n\rho_2} \mathbf{1} \left(h^{(i,j)} + p^{(i,j)} - \mathbf{1}^\top (g^{(i,j)}(x^{(i,j)}) - \lambda^{(i,j)}) \right) \quad (7b)
\end{aligned}$$

$$x_1^{(i)} \leftarrow \sum_j (x_1^{(i,j)} - \nu^{(i,j)}) \quad (7c)$$

$$h^{(i,j)} \leftarrow \Pi_{(-\infty, c(i)]} (\mathbf{1}^\top z^{(i,j)} - p^{(i,j)}) \quad (7d)$$

$$s^{(j)} \leftarrow \Pi_{[0, \infty)} \left(\sum_{i=1}^m x^{(i,j)} - y^{(j)} - \mu^{(j)} \right) \quad (7e)$$

$$\mu^{(j)} \leftarrow \mu^{(j)} + s^{(j)} - \sum_i x^{(i,j)} + y^{(j)} \quad (7f)$$

$$\lambda^{(i,j)} \leftarrow \lambda^{(i,j)} + z^{(i,j)} - g^{(i,j)}(x^{(i,j)}) \quad (7g)$$

$$\nu^{(i,j)} \leftarrow \nu^{(i,j)} + x_1^{(i)} - x_1^{(i,j)} \quad (7h)$$

$$p^{(i,j)} \leftarrow p^{(i,j)} + h^{(i,j)} - \mathbf{1}^\top z^{(i,j)} \quad (7i)$$

where $\Pi_{\mathbb{I}}(x)$ for scalar x denotes projection onto the interval \mathbb{I} , while for vector x , $\Pi_{\mathbb{I}}(x)$ is the vector obtained by projecting each element of x onto \mathbb{I} , and $\delta_{k,1} = 1$ if $k = 1$, $\delta_{k,1} = 0$ if $k \neq 1$. Note that the update (7a) can be written equivalently as

$$\begin{aligned}
x_k^{(i,j)} \leftarrow \Pi_{[\underline{x}_k^{(i)}, \bar{x}_k^{(i)}]} \left(\underset{x_k^{(i,j)}}{\operatorname{argmin}} \left\{ \frac{1}{q} f_k^{(i,j)}(x_k^{(i,j)}) \right. \right. \\
+ \frac{\rho_1}{2} (\theta_k^{(i,j)} - g_k^{(i,j)}(x_k^{(i,j)}))^2 + \frac{\rho_3}{2} (\phi_k^{(i,j)} - x_k^{(i,j)})^2 \\
\left. \left. + \delta_{k,1} \frac{\rho_4}{2} (x_1^{(i)} - x_1^{(i,j)} + \nu^{(i,j)})^2 \right\} \right).
\end{aligned}$$

with $\theta_k^{(i,j)} = z_k^{(i,j)} + \lambda_k^{(i)}$ and $\phi_k^{(i,j)} = s_k^{(j)} - \sum_{l \neq i} x_k^{(l,j)} + y_k^{(j)} + \mu_k^{(j)}$. The iteration terminates when the primal and dual residuals r and σ defined by

$$\begin{aligned}
r &= \max \left\{ \max_j \|s^{(j)} - \sum_i x^{(i,j)} + y^{(j)}\|_\infty, \right. \\
&\quad \max_{i,j} \|z^{(i,j)} - g^{(i,j)}(x^{(i,j)})\|_\infty, \\
&\quad \left. \max_{i,j} |h^{(i,j)} - \mathbf{1}^\top z^{(i,j)}|, \max_{i,j} |x_1^{(i,j)} - x_1^{(i)}| \right\} \\
\sigma &= \max \left\{ \rho_1 \max_i \|z^{(i)} - \tilde{z}^{(i)}\|_\infty, \rho_2 \max_{i,j} |h^{(i,j)} - \tilde{h}^{(i,j)}|, \right. \\
&\quad \left. \rho_3 \max_{i,j} \|s^{(j)} - \tilde{s}^{(j)} - \sum_i (x^{(i,j)} - \tilde{x}^{(i,j)})\|_\infty \right\}
\end{aligned}$$

(where $[(\cdot)]_0$ denotes the value of an iterate (\cdot) at the previous iteration), fall below predefined thresholds, \bar{r} and $\bar{\sigma}$.

REFERENCES

- [1] M. Koot, J. Kessels, B. de Jager, W. Heemels, P. van den Bosch, and M. Steinbuch, "Energy management strategies for vehicular electric power systems," *IEEE Transactions on Vehicular Technology*, vol. 54, no. 3, pp. 771–782, 2005.
- [2] J. Buerger, S. East, and M. Cannon, "Fast dual-loop nonlinear receding horizon control for energy management in hybrid electric vehicles," *IEEE Trans Control Systems Technology*, vol. –, pp. 1–11, 2018.
- [3] P. van den Bosch and F. Lootsma, "Scheduling of power generation via large-scale nonlinear optimization," *Journal of Optimization Theory and Applications*, vol. 55, no. 2, pp. 313–326, 1987.
- [4] H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
- [5] M. Patriksson, "A survey on the continuous nonlinear resource allocation problem," *EJOR*, vol. 185, no. 1, pp. 1–46, 2008.
- [6] M. Campi and S. Garatti, "The exact feasibility of randomized solutions of uncertain convex programs," *SIAM Journal on Optimization*, vol. 19, pp. 1211–1230, 2008.
- [7] J. Eckstein, "Parallel alternating direction multiplier decomposition of convex programs," *Journal of Optimization Theory and Applications*, vol. 80, pp. 39–62, 1994.
- [8] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 11, pp. 1–122, 2011.
- [9] G. Reis, F. Zeilfelder, M. Hering-Bertram, G. Farin, and H. Hagen, "High-quality rendering of quartic spline surfaces on the GPU," *IEEE Transactions on visualization and computer graphics*, vol. 14, no. 5, pp. 1126–1139, 2008.
- [10] R. Klotek and J. Porter-Sobieraj, "Solving systems of polynomial equations on a GPU," *2012 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pp. 539–544, 2012.
- [11] *CUDA C Programming Guide v9.2.148*, NVIDIA Corporation, 2018. [Online]. Available: <https://docs.nvidia.com/cuda/#>
- [12] B. Egardt, N. Murgovski, M. Pourabdollah, and L. Johannesson Mardh, "Electromobility studies based on convex optimization: Design and control issues regarding vehicle electrification," *IEEE Control Systems*, vol. 34, no. 2, pp. 32–49, 2014.
- [13] S. Hadj-Said, G. Colin, A. Ketfi-Cherif, and Y. Chamaillard, "Convex optimization for energy management of parallel hybrid electric vehicles," *IFAC-PapersOnLine*, vol. 49, no. 11, pp. 271–276, 2016.
- [14] G. Calafiore, "Random convex programs," *SIAM Journal on Optimization*, vol. 20, no. 6, pp. 3427–3464, 2010.
- [15] D. Dummit and R. Foote, *Abstract Algebra*. John Wiley, 2004.
- [16] M. Josevski, A. Katrinik, and D. Abel, "Scenario MPC for fuel economy optimization of hybrid electric powertrains on real-world driving cycles," *Proc American Control Conf*, pp. 5629–5635, 2017.
- [17] G. Ripaccioli, D. Bernardini, S. Di Cairano, A. Bemporad, and I. V. Kolmanovskiy, "A stochastic model predictive control approach for series hybrid electric vehicle power management," *Proc American Control Conf*, pp. 5844–5849, 2010.
- [18] S. East and M. Cannon, "ADMM for MPC with state and input constraints, and input nonlinearity," in *Proc American Control Conf*, 2018, pp. 4514–4519.