Time-of-Use Pump Scheduling for Flow Transmission*

Yuzhen Feng

Abstract

We study time-of-use pump scheduling to deliver a required volume using a finite set of pump combinations with empirical flow–power performance, subject to per-shift caps on pump switches. We prove a structural theorem: partitioning the horizon into maximal intervals with constant tariff and shift (atoms), there always exists an optimal schedule with at most one internal switch in at most one atom, with all other switches at atom boundaries. This insight yields an exact mixed-integer linear program with atom-level variables and without time discretization.

1 Introduction

Pumping systems are workhorses of continuous-process industries such as crude-oil gathering, long-distance pipeline transport, water distribution, and refinery transfer operations. Operators must schedule pumps over a planning horizon under time-varying electricity prices, typically time-of-use (TOU) with peak and off-peak rates (e.g., Table 1). Appropriate schedules can therefore reduce energy cost while meeting delivery targets.

Table 1: TOU electricity tariffs.

Start (h)	End (h)	Tier	Price (CNY/kWh)
0	8	Off-peak	0.21
8	10	Shoulder	0.55
10	14	Peak	0.94
14	19	Shoulder	0.55
19	24	Peak	0.94

We consider constant-speed pumps operated singly or in multi-pump combinations (combos). Because of hydraulic interactions, combo throughput is not additive; instead, each combo is represented by a measured pair of volumetric flow and electric power, giving a finite set of feasible operating points (e.g., Table 2). Exactly one combo is active at any time; outages can be modeled by a zero-flow, zero-power idle combo. In addition to energy

*This work originated from Sinopec's First AI Innovation Competition.

cost, real plants limit how often crews can change the active combo, as frequent cycling increases wear and labor requirements. We can capture this with a per-shift switch cap: within each crew shift, the number of changes from one combo to another is bounded. Our analysis and algorithms extend directly to a per-switch penalty.

Table 2: Measured performance of pump combos.

Combination	Flow (m^3/h)	Power (kW)
1#	1,055	1,102
2#	1,113	1,120
1#&2#	1,880	2,352

Given (i) a TOU price profile, (ii) a finite set of combos with measured performance, (iii) a required delivery volume, and (iv) per-shift caps on the number of switches, the scheduler must choose a time sequence of combos that meets the volume at minimum electricity cost. A straightforward approach discretizes time into fine slots and yields a time-indexed integer linear program with on/off binary variables for each combo and slot. Such models are easy to formulate but face a trade-off between accuracy and scalability: finer discretization approximates continuous time better, yet the number of binary variables grows with the number of slots and quickly strains solvers over realistic horizons.

Our main contribution is a structural theorem for TOU pump scheduling. Partitioning the horizon into maximal intervals with constant tariff and shift (atoms), we show that there exists an optimal schedule with at most one internal switch in at most one atom, with all remaining switches placed at atom boundaries. This insight can yield exact formulations that avoid fine time discretization. Building on this result, we formulate a mixed integer linear program (MILP) that operates at the atom level and explicitly handles the unique possible internal switch.

2 Problem Setting

We consider a task starting at time $t = 0$ to deliver a volume $V > 0$. Let $\mathcal{K}$ be a finite set of pump combos, where each $k \in \mathcal{K}$ has constant flow $v_k \geq 0$ and power $q_k \geq 0$. The TOU tariff is a right-continuous, piecewise-constant function $c : \mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$, and crew shifts are described by a right-continuous, piecewise-constant mapping $s : \mathbb{R}_{\geq 0} \to \mathcal{S}$. Each shift has a cap $W \in \mathbb{Z}_{\geq 0}$ on the number of switches attributed to it. Let $0 = t_0 < t_1 < \cdots < t_A$ denote all time points at which either c or s changes. Define atoms $a \in \mathcal{A} = \{0, \ldots, A-1\}$ as the intervals $[t_a, t_{a+1})$, with length $l_a = t_{a+1} - t_a$, price $c_a = c(t_a)$, and shift label $s_a = s(t_a)$. By construction, both c and s are constant on each atom.

A schedule is a right-continuous, piecewise-constant mapping $k : [0, T^*) \to \mathcal{K}$ for some

completion time $T^* > 0$. A switch occurs at any $t \in (0, T^*)$ with $k(t^-) \neq k(t^+)$. The schedule must deliver exactly

$$V = \int_0^{T^*} v_{k(t)} \mathrm{d}t$$

subject to the per-shift switch caps W, and the objective is to minimize total energy cost

$$\int_0^{T^*} c(t) q_{k(t)} \mathrm{d}t.$$

Illustration. Figure 1 shows an optimal schedule, starting at midnight, that uses the combos in Table 2 under the tariff profile in Table 1 to meet a volume target of 40,000m^3. Although single pump 2# achieves the lowest specific energy (kWh/m^3), the schedule front-loads delivery by running combo 1#&2# during off-peak hours, thereby reducing the need to operate during subsequent peak periods.

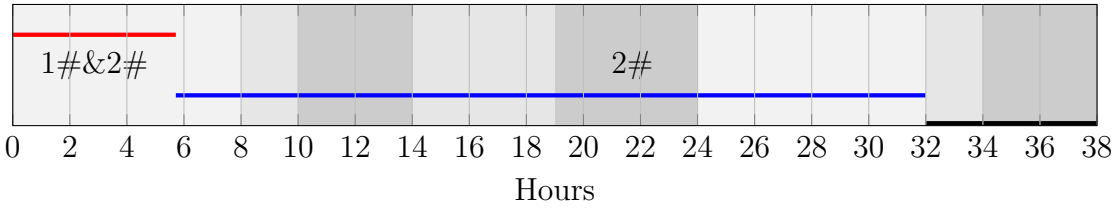

Figure 1: Example schedule; background shading indicates the TOU tariff tier.

3 Structural Properties

We now work with the continuous, event-based timeline from Section 2 and derive structural properties of optimal schedules. For each atom $a \in \mathcal{A}$, define the left- and right-limit combos

$$k_a^{\leftarrow} := k(t_a^+), \quad k_a^{\rightarrow} := k(t_{a+1}^-).$$

Thus $k_a^{\leftarrow}$ is the combo active immediately after entering atom a, and $k_a^{\rightarrow}$ is the combo active immediately before leaving it. When $k_a^{\leftarrow} \neq k_a^{\rightarrow}$, at least one switch occurs within the atom.

3.1 Within-Atom Combo Assignment

Fix an atom a that is used for its full length l_a. Any within-atom pattern can be summarized by time fractions $\{x_{a,k}\}_{k\in\mathcal{K}}$ of l_a satisfying

$$\sum_{k\in\mathcal{K}} x_{a,k} = 1, \tag{1a}$$

$$\sum_{k\in\mathcal{K}} l_a v_k x_{a,k} = V_a, \tag{1b}$$

$$x_{a,k} \geq 0,\ \forall k \in \mathcal{K}. \tag{1c}$$

For fixed V_a, minimizing per-atom energy cost $C_a = c_a \sum_k l_a q_k x_{a,k}$ is a linear program (LP).

Lemma 1. *For fixed a and $V_a \in [0, \max_k v_k l_a]$, there exists an optimal solution to (1) with at most two positive $x_{a,k}$; that is, some optimal schedule uses at most two distinct combos.*

Proof. The feasible region is a polyhedron defined by two equalities and nonnegativity, so its extreme points have at most two positive components. □

Within atom a, write the block pattern on $[t_a, t_{a+1})$ as

$$\left[k^{(1)} \mid k^{(2)} \mid \cdots \mid k^{(m)}\right],$$

where $m \geq 1$, $k^{(1)} = k_a^{\leftarrow}$, $k^{(m)} = k_a^{\rightarrow}$, and each "|" denotes an internal switch within atom a.

Proposition 1. *There exists an optimal schedule in which every fully used atom has at most one internal switch (i.e., $m \leq 2$ on each such atom).*

Proof. Fix a globally optimal schedule and a fully used atom a with m blocks. If $m \geq 3$, Lemma 1 provides an equal- or lower-cost allocation using at most two combos $\{i, j\}$ for the same V_a. Implementing it as two contiguous blocks $[i|j]$ reduces the internal switch count in a to at most one. This transformation does not increase the number of switches attributed to any shift, so feasibility and optimality are preserved. If $m \leq 2$, the claim already holds. □

Proposition 2. *There exists an optimal schedule in which the partially used (last) atom, if any, has no internal switch; it uses a single combo.*

Proof. On the partially used last atom, the normalization constraint (1a) is slack, so the within-atom LP has one equality and nonnegativity. An extreme-point optimum therefore has exactly one positive component, corresponding to a single combo. □

Proposition 1 allows us to restrict attention to schedules with at most one internal switch on any fully used atom, and Proposition 2 removes internal switch on a partially used terminal atom.

3.2 Across-Atom Volume Allocation

Given the within-atom structure, we now characterize feasible (V_a, C_a) pairs for each atom.

Fully used atom with no internal switch. If $k_a^{\leftarrow} = k_a^{\rightarrow} = k_a$, then

$$V_a = l_a v_{k_a}, \quad C_a = c_a l_a q_{k_a}.$$

Fully used atom with one internal switch. If $k_a^{\leftarrow} \neq k_a^{\rightarrow}$, let $\{k_a^{\min}, k_a^{\max}\} = \{k_a^{\leftarrow}, k_a^{\rightarrow}\}$ be ordered so that $v_{k_a^{\min}} \leq v_{k_a^{\max}}$. With at most one internal change, any

$$V_a \in l_a \left[v_{k_a^{\min}}, v_{k_a^{\max}}\right]$$

is attainable by splitting time shares. Writing $\theta_a(V_a) \in [0,1]$ for the share on $k_a^{\min}$,

$$V_a = l_a \left[\theta_a(V_a) v_{k_a^{\min}} + (1 - \theta_a(V_a))\, v_{k_a^{\max}}\right],$$

so $C_a(V_a)$ is affine on this interval:

$$C_a(V_a) = c_a l_a \left[\theta_a(V_a) q_{k_a^{\min}} + (1 - \theta_a(V_a))\, q_{k_a^{\max}}\right].$$

Terminal partial atom. Let a° be the last used atom. By Proposition 2 we use a single-combo segment with left endpoint $k_{a^\circ}^{\leftarrow}$, hence

$$V_{a^\circ} \in l_{a^\circ} \left[0, v_{k_{a^\circ}^{\leftarrow}}\right], \quad C_{a^\circ}(V_{a^\circ}) = c_{a^\circ} \frac{V_{a^\circ}}{v_{k_{a^\circ}^{\leftarrow}}} q_{k_{a^\circ}^{\leftarrow}}.$$

Fix the endpoint pairs $(k_a^{\leftarrow}, k_a^{\rightarrow})_{a \in \mathcal{A}}$ and the switch attributions. The remaining decision is the across-atom volume allocation V_a. The problem is to minimize

$$\sum_{a \in \mathcal{A}} C_a(V_a) \tag{2}$$

subject to the total volume and per-atom feasibility:

$$\sum_{a\in\mathcal{A}} V_a = V, \tag{3a}$$

$$V_a \in \mathcal{V}_a := \begin{cases} \{l_a v_{k_a}\}, & \text{if } k_a^{\leftarrow} = k_a^{\rightarrow} = k_a \text{ and } a \text{ is fully used,} \\ l_a \left[v_{k_a^{\min}}, v_{k_a^{\max}}\right], & \text{if } k_a^{\leftarrow} \neq k_a^{\rightarrow} \text{ and } a \text{ is fully used,} \\ l_a \left[0, v_{k_a^{\leftarrow}}\right], & \text{if } a \text{ is the partially used terminal atom.} \end{cases} \tag{3b}$$

Lemma 2. *There exists an optimal solution $(V_a^*)_{a\in\mathcal{A}}$ to (3) with at most one index $a^\star$ such that $V_{a^\star}^*$ lies in the interior of $\mathcal{V}_{a^\star}$; all other V_a^* are at their bounds (or fixed).*

Proof. Problem (3) is an LP with one coupling equality and box constraints; extreme points have at most one variable not at a bound. □

Combining the above: from any optimal schedule we can construct another optimal schedule in which every fully used atom has at most one internal switch, any partially used last atom is single-combo, and once the endpoint pairs are fixed, at most one atom is split internally.

Theorem 1. *There exists an optimal schedule with the following structure:*

1. *If the last used atom is partial, then no atom is split internally; all switches occur at atom boundaries, and the last atom is single-combo.*

2. *If the schedule ends at an atom boundary (no partial last atom), then at most one atom is split internally; all other switches occur at atom boundaries.*

4 Atom-Level MILP

Leveraging the structural properties in Section 3, we formulate an atom-level MILP that allows at most one internal switch over the entire horizon. Let $\mathcal{P} = \{0, \ldots, P-1\}$ index the shifts. For each shift $p \in \mathcal{P}$, let $\mathcal{A}_p \subset \mathcal{A}$ be its atoms, let $a_p^{\triangleright}$ be the first atom in shift p, and define $\mathcal{B}_p := \mathcal{A}_p \setminus \{a_p^{\triangleright}\}$ as the atoms whose left boundary lies strictly inside shift p.

Decision variables. The MILP uses the decision variables listed in Table 3.

The binary variables η_p and the auxiliary variables α_p implement flexible attribution of switches that occur exactly at a shift boundary. In the continuous-time model, such a boundary switch can be counted either in the preceding shift or in the following one, as long as the per-shift cap W is respected. In the MILP, η_p selects whether the switch at the start of shift p is attributed to shift $p-1$ or to shift p, and α_p is used to linearize the product $\eta_p \sigma_{a_p^{\triangleright}}$ via McCormick envelopes.

Table 3: Decision variables of the atom-level MILP.

Variable	Domain	Description
y_a	$\{0,1\}$	1 if atom $a \in \mathcal{A}$ is used
$\bar{y}_a$	$\{0,1\}$	1 if atom $a \in \mathcal{A}$ is the terminal atom
ρ_a	$[0,1]$	Fraction of atom $a \in \mathcal{A}$ used if terminal
z_a	$\{0,1\}$	1 if atom $a \in \mathcal{A}$ is internally split once
$x^{\leftarrow}_{a,k}$	$\{0,1\}$	1 if combo $k \in \mathcal{K}$ is active at t^+_a, $a \in \mathcal{A}$
$x^{\rightarrow}_{a,k}$	$\{0,1\}$	1 if combo $k \in \mathcal{K}$ is active at t^-_{a+1}, $a \in \mathcal{A}$
$t^{\leftarrow}_{a,k}$	$\mathbb{R}_{\geq 0}$	Run time of combo $k \in \mathcal{K}$ in the left block of atom $a \in \mathcal{A}$
$t^{\rightarrow}_{a,k}$	$\mathbb{R}_{\geq 0}$	Run time of combo $k \in \mathcal{K}$ in the right block of atom $a \in \mathcal{A}$
σ_a	$\{0,1\}$	1 if there is a switch at boundary t_a, $a \in \mathcal{A}$; $\sigma_0 = 0$
η_p	$\{0,1\}$	1 if the switch at the start of shift $p \in \mathcal{P}$ is attributed to shift $p-1$
α_p	$[0,1]$	Auxiliary variable for $\alpha_p = \eta_p \sigma_{a^{\triangleright}_p}$, $p \in \mathcal{P}$

Objective.

$$\min \quad \sum_{a \in \mathcal{A}} c_a \sum_{k \in \mathcal{K}} q_k \big(t^{\leftarrow}_{a,k} + t^{\rightarrow}_{a,k}\big). \tag{4}$$

Constraints.

$$y_0 = 1, \tag{5a}$$

$$y_a \leq \min\{y_{a-1}, 1 - \bar{y}_{a-1}\}, \ \forall a \in \mathcal{A} \setminus \{0\}, \tag{5b}$$

$$\sum_{a \in \mathcal{A}} \bar{y}_a = 1, \tag{5c}$$

$$\rho_a \leq \bar{y}_a \leq y_a, \ \forall a \in \mathcal{A}. \tag{5d}$$

$$\sum_{a \in \mathcal{A}} z_a \leq 1. \tag{6}$$

$$\sum_{k \in \mathcal{K}} x^{\leftarrow}_{a,k} = \sum_{k \in \mathcal{K}} x^{\rightarrow}_{a,k} = y_a, \ \forall a \in \mathcal{A}, \tag{7a}$$

$$x^{\rightarrow}_{a,k} - x^{\leftarrow}_{a,k} \leq z_a, \ \forall k \in \mathcal{K}, \ a \in \mathcal{A}. \tag{7b}$$

$$\sum_{k \in \mathcal{K}} \big(t^{\leftarrow}_{a,k} + t^{\rightarrow}_{a,k}\big) = l_a(y_a - \bar{y}_a) + l_a \rho_a, \ \forall a \in \mathcal{A}, \tag{8a}$$

$$t^{\leftarrow}_{a,k} \leq l_a x^{\leftarrow}_{a,k}, \ t^{\rightarrow}_{a,k} \leq l_a x^{\rightarrow}_{a,k}, \ \forall a \in \mathcal{A}, \ k \in \mathcal{K}, \tag{8b}$$

$$\sum_{k \in \mathcal{K}} t^{\rightarrow}_{a,k} \leq l_a z_a, \ \forall a \in \mathcal{A}. \tag{8c}$$

$$\sum_{a\in\mathcal{A}}\sum_{k\in\mathcal{K}} v_k\big(t^{\leftarrow}_{a,k} + t^{\rightarrow}_{a,k}\big) = V. \tag{9}$$

$$\sigma_0 = 0, \tag{10a}$$

$$x^{\leftarrow}_{a,k} - x^{\rightarrow}_{a-1,k} \le \sigma_a \le y_a,\ \forall k \in \mathcal{K},\ a \in \mathcal{A}\setminus\{0\}. \tag{10b}$$

$$\alpha_p \le \eta_p,\ \forall p \in \mathcal{P}, \tag{11a}$$

$$\alpha_p \le \sigma_{a^{\triangleright}_p},\ \forall p \in \mathcal{P}, \tag{11b}$$

$$\alpha_p \ge \eta_p + \sigma_{a^{\triangleright}_p} - 1,\ \forall p \in \mathcal{P}, \tag{11c}$$

$$\alpha_p \ge 0,\ \forall p \in \mathcal{P}. \tag{11d}$$

$$\sum_{a\in\mathcal{A}_0} z_a + \sum_{a\in\mathcal{B}_0} \sigma_a + \alpha_1 \le W, \tag{12a}$$

$$\sum_{a\in\mathcal{A}_p} z_a + \sum_{a\in\mathcal{B}_p} \sigma_a + \sigma_{a^{\triangleright}_p} - \alpha_p + \alpha_{p+1} \le W,\ p = 1,\dots,P-2, \tag{12b}$$

$$\sum_{a\in\mathcal{A}_{P-1}} z_a + \sum_{a\in\mathcal{B}_{P-1}} \sigma_a + \sigma_{a^{\triangleright}_{P-1}} - \alpha_{P-1} \le W. \tag{12c}$$

The objective (4) minimizes total cost. Constraints (5) enforce contiguity, select a unique terminal atom, and activate the terminal fraction ρ_a only there. Constraint (6) imposes the global "at most one internal split" property. Constraints (7) choose the endpoint combos for each atom and force them to coincide unless $z_a = 1$. Constraints (8) link endpoint choices to run times and ensure that fully used atoms contribute l_a hours while the terminal atom contributes $\rho_a l_a$. Constraint (9) enforces the total volume V. Constraints (10) detect switches at atom boundaries. Constraints (11) linearize the flexible attribution of the switch at each shift boundary. Finally, constraints (12) count internal splits and boundary switches to enforce the per-shift cap W.