Challenge: Modelling Unit Commitment as a Planning Problem

Joshua Campion, Chris Dent
School of Engineering and Computing Sciences
University of Durham, UK
email: {*j.r.campion,chris.dent*}@*durham.ac.uk*

Maria Fox, Derek Long, Daniele Magazzeni
Department of Informatics
King's College London UK
email:{*firstname.lastname*}@*kcl.ac.uk*

Abstract

Unit Commitment is a fundamental problem in power systems engineering, deciding which generating units to switch on, and when to switch them on, in order to efficiently meet anticipated demand. It has traditionally been solved as a Mixed Integer Programming (MIP) problem but upcoming changes to the power system drastically increase the MIP solution time. In this paper, we discuss the benefits that using planning may have over the established methods. We provide a formal description of Unit Commitment, and we present its formulation as MIP and as a planning problem. This is a novel and interesting application area for planning, with features that make the domain challenging for current planners.

Introduction

Unit Commitment (UC) is deciding which generating units in a power system to switch on at what time. Economic Dispatch is the subsequent problem of, given which generating units are currently switched on, setting the output levels of each unit (Wollenberg and Wood 1996). They are discussed in tandem as dispatch models are embedded within UC models in order to analyse which combination of units would be best. Changes to the bulk power system imply there must be changes to existing methods.

Electricity generation accounted for ∼24% of carbon emissions in 2006 and one point projection of electricity consumption expects an increase of 77% from 2006 levels by 2030 (Agency 2012) but EU governments demand a reduction to 20% below 1990 levels by 2020 as part of the 20-20-20 climate change initiative (Commission 2012). This is clearly not possible without a dramatic decarbonisation of the electricity network. Wind generation in the UK alone has grown to over 7 GW at time of writing (see (RenewableUK 2012) for latest statistics) and is projected to increase to 30-45GW by 2030 (on Climate Change 2012; Consulting 2009). The variability of wind and other renewables is in stark contrast to entirely controllable traditional power systems (Ilic 2007). To reduce the curtailment of wind power system operators must predict upcoming wind power output and reduce scheduled thermal output accordingly, adding stochasticity to the UC problem.

Controllability of traditional generation (aside from unit failures, probabilities of which can be statistically aggregated with confidence) allowed construction of a deterministic optimisation problem, formed as a MIP and solved, using for example, Branch and Bound or Lagrangian Relaxation (Streiffert, Philbrick, and Ott 2005). The aforementioned changes mean traditional formulations of the problem should be reconsidered to take the variability and non-controllability of renewable resources into account.[1]

One commonly proposed solution (see for example (Barth et al. 2006; Meibom et al. 2007; 2011; Tuohy et al. 2009; Sturt and Strbac 2012)) is to combine Rolling Planning and Stochastic UC. These methods involve resolving the MIP model for different net demand (demand minus wind) profiles. This is computationally very intensive when using a fine discretisation of wind forecast distributions and a new faster solution method would be beneficial to the current power systems community.

In this paper we present a new challenging domain for planning and we discuss the benefits that using planning may have over the established methods. A formal description of the Unit Commitment problem is presented, followed by its formulation as MIP and as a planning problem. Features of the domain that make it challenging for current planners are also discussed.

Benefits of AI Planning

We believe there is potential for a faster solution of comparable quality using AI Planning. One downside of the MIP is its fixed discretisation of the time points at which changes can occur. This leads to the binary on / off variables being set at the same value for long periods of time, and solution time is spent reasoning this. A planning formulation differs from all treatments of Unit Commitment in the literature (MIP based formulations) in that it is formulated as actions rather than states at set time points.

One major benefit of a planning formulation over a MIP formulation is that the number of actions required to meet

[1]Note that traditional UC formulations could remain unchanged if larger reserve requirements were imposed, however this would increases cost.

the load should be much less than the number of time points times the number of units. The planning model we propose, when guided correctly, could leave the system in steady state and only reason on any necessary actions to meet the changed demand when sufficient changes occur.

A secondary benefit of solving Unit Commitment with AI Planning is that by reasoning in continuous time the objective function of the Unit Commitment schedule can be futher optimised. Furthermore, heuristic search can allow scaling to big problems in order to deal with realistic instances where many units are considered. This is motivated by this problem's similarity to that in (Fox, Long, and Magazzeni 2011) on plan-based management of battery load. In this work, a very effective heuristic is embedded in the UPMurphi planner (Della Penna et al. 2009), which found extremely high-quality solutions for the efficient scheduling of multiple batteries.

In this paper we present a deterministic model of the problem, where we assume to know the demand to service in advance. In a more realistic scenario, a probabilistic distribution characterising typical demands can be learned from the experience, and an approach based on classification could be used, as proposed by (Fox, Long, and Magazzeni 2012).

First, we provide the main components of the MIP formulation of the Unit Commitment. As a second contribution of this paper, we present the Unit Commitment as a planning problem as it presents a mix of constraints and features which are challenging for current planners. In particular, costs need to be computed continuously as a function of the output of each unit being used. The invariant condition of the total output being greater than the demand must be satisfied, subject to changes in the deman modelled as a set of Timed Initial Fluents. To the best of our knowledge, no planner is currently able to solve this problem, because of the invariant conditions dependent of Timed Initial Fluents.

The planner COLIN (Coles et al. 2012) has been used to tackle domains with continuous linear change but this problem requires actions in response to Timed Initial Fluents. This has not been modelled before and is discussed in more detail prior to the model description.

Problem Formulation

We formalise the problem here to ensure both models use the same set of assumptions.

- n generating units (indexed with u), each with the characteristics given below, are to serve a deterministic, piecewise constant, demand over a horizon of length T.
  - Minimum and maximum stable generation levels $G_{u,\text{min}}$, $G_{u,\text{max}}$ outside of which the generating unit u cannot operate.
  - A fixed start up time $T_{u,\text{start}}$. From on offline state, this is the amount of time it takes for the generating unit u to come online and output power. During this time no power is output, and after this time the output is the minimum stable generation level.
  - A fixed switch off time $T_{u,\text{off}}$. From an online state, this is the amount of time it takes for the generating unit u to come offline. The generating unit u must be outputting its minimum stable generation level before it can be switched off. Once switching off begins no power is output.
  - A fixed minimum run and minimum off time, $T_{u,\text{min on}}, T_{u,\text{min off}}$. Once a unit is comes online (offline) it must remain online (offline) for at least $T_{u,\text{min on}}$ ($T_{u,\text{min off}}$).
  - The maximum increase (decrease) in output of a generating unit u is equal to $R_{u,+(-)}$ MW / min. This is known as the ramp rate.
  - A start up cost $C_{u,\text{start}}$ (€). The cost of switching on a generating unit u.
  - A no-load cost $C_{u,\text{no-load}}$ (€/ min). The cost of having the generating unit u online regardless of output.
  - A marginal running cost $C_{u,\text{marginal}}$ (€/ min). The per MW cost of output from the generating unit u.
- The total output from all generating units must be greater than the demand at all times.[2]
- The total demand should be served at the lowest cost.
- The system will have an initial configuration given by the end state of the previous planning horizon. Switching a unit on or off near the end of the horizon will mean it may not serve the full $T_{u,\text{min on}}$ or $T_{u,\text{min off}}$, creating starting parameters $T_{u,\text{initial on}}$ and $T_{u,\text{initial off}}$ which will be fixed for a given problem instance.

MIP Model

The above problem can be easily formulated as a MIP problem. Let $\mathcal{U}$ be the set of generating units indexed by u, $i = 1, \ldots, m$ index the time periods over which the demand is constant and $\Delta t = T/m$ be the length of those intervals. All fixed parameters are as detailed above. The decision variables are then: $o(u, i)$ a binary flag for whether or not a generating unit u is online during a given period i, $g(u, i)$, the exact output of the generating unit u during a given period, and $c(u, i)$ an extra variable tracking cost of start ups.

The objective function should minimise total operation cost, and is given by

$$f := \sum_{i=1}^{m} \sum_{u \in \mathcal{U}} \Big[g(u,i) \cdot C_{u,\text{marginal}} \cdot \Delta t + o(u,i) \cdot C_{u,\text{no-load}} \cdot \Delta t + c(u,i) \Big] \tag{1}$$

There have been many different models proposed in the literature, (Ostrowski, Anjos, and Vannelli 2012) compare two models, one with the same decision variables described here and one with extra binary variables. The extra binaries allow for the construction of tighter constraints which can reduce the solution time. It is important that we use the latest MIP models to give a representative comparison of solution methods.

[2]The amount by which supply exceeds demand is known as the spinning reserve. Many different reserve strategies have been presented in the literature, it is not within the scope of this paper to discuss the merits of these.

All constraints from the above formulation can be imposed as linear constraints on the above decision variables using common formulations such as those in (Carrión and Arroyo 2006).

AI Planning Model

Previous Work

This problem presents a mix of features which have not been addressed together before. In fact, there is continuous change involving the cost computation, which depends on the output of each unit being used. The demand to be services is modelled as a sequence of Timed Initial Fluents, which influence an invariant condition that must be satisfied during the plan to ensure the total output is greater than the demand.

This is similar to both (Fox, Long, and Magazzeni 2011), where multiple batteries are tasked to serve a deterministic load profile over a given horizon, and (Bell et al. 2008), where actions are taken to ensure voltage levels remain inside a predefined range following a changing demand profile, but differs the in following ways.

This problem differs from (Fox, Long, and Magazzeni 2011) in the need for concurrent actions and coordination between multiple objects, the generating units. There, only one battery acts at a time and the reasoning required is to switch which battery is serving the current load. The problem also differs in size. The battery problem demonstrates performance for up to 8 objects, whereas realistic instances of Unit Commitment involve up to 50 units. This is a huge increase in the number of 'interesting' choices at each decision point and scaling well will be a challenge.

In (Bell et al. 2008) only one control action is performed at any one time as only one substation is to be controlled. In Unit Commitment all operating units must be considered simultaneously to ensure demand is continually served. Also the domain there is non-temporal so there is no need for continuous reasoning about the changing demand.

UC Planning Domain

In the following we describe the planning domain, highlighting how the constraints and features discussed above are modelled in PDDL2.1(Fox and Long 2003).

The invariant constraint of the total supply being greater than demand is modelled as an `over all` condition of an *envelope* action modelling the whole task of servicing the complete load for the 24 hours, as shown in Figure 1.

The predicate `complete` which appears in the goal is added only once this actions is complete. This forces the planner to run this action, forcing ensuring that `(>= (supply) (demand))` is always satisfied. Note that this must be both an over all and boundary constraint to ensure the demand is met at all times. On the other hand, the start effect adds the `running` predicate, which is a start condition of all other actions. This forces the planner to run `serve-load` as the first action. As we already said, the demand profile is altered by Timed Initial Fluents which change the fluent `demand` at specified time points.

```
(:durative-action serve-load
 :parameters ()
 :duration (= ?duration 1440)
 :condition ( and
     (at start (>= (supply) (demand)))
     (over all (>= (supply) (demand)))
     (at end (>= (supply) (demand))))
 :effect (and (at start (running))
         (at end (complete))))
```

Figure 1: Envelope modelling the invariant condition of servicing the demand.

```
(:durative-action switch-on
 :parameters (?u - unit)
 :duration (= ?duration (switch-on-time ?u))
 :condition (and
       (at start (running))
       (at start (off ?u))
       (at start (canSwitchOn ?u)))
 :effect (and
       (at start (not (canSwitchOn ?u)))
       (at start (not (off ?u)))
       (at start (increase (totalCost)
                 (startup-cost ?u)))
       (at end (on ?u))
       (at end (assign (output ?u)
               (min-output  ?u)))))
```

Figure 2: The durative action `switch-on`.

The n generating units will be represented by n objects. The units' properties will be modelled as follows:

Minimum on / off time: Two predicates `canSwitchOn` and `canSwitchOff` determine whether switch on and off actions can be applied. These allow for the implementation of minimum on / off time by controlling whether or not the planner can alter a units on / off state. Initial on / off times can be set by combining these predicates with Timed Initial Literals.

Start up time: for each unit, the fixed start up time is modelled in the `switch-on` action shown in Figure 2. The output of the single unit is set by this method, but without updating the total supply, which is instead updated in the action `generate` which keeps track of the continuous costs. A similar `switch-off` action coupled with a `canSwitcOff` predicate is used to set the output of the generating unit u to 0 and fix the switch off time.

Continuous cost: in order to continuously track the marginal costs, three actions are used: `generate`, `ramp-up` and `ramp-down`. The action `generate`, shown in Figure 3, is responsible for updating the total supply, incrementing it when a unit is on, and decrementing it when the unit is no longer used. The continuous change of PDDL2.1 is used to continuously update the total cost, by taking into account the marginal cost and the current output of the unit.

```
(:durative-action generate
 :parameters (?u - unit)
 :duration (<= ?duration 1440)
 :condition (and (at start (running))
                 (over all (on ?u))
 :effect (and  (increase (totalCost)
 (* #t ( * (output ?u) (marginal-cost ?u))))
 (at start (increase (supply) (output ?u)))
 (at end (decrease (supply) (output ?u))))))
```

Figure 3: Durative action `generate`

```
(:durative-action ramp-up
 :parameters (?u - unit)
 :duration (<= ?duration
       (/ (- (max-output ?u) (output ?u))
          (up-ramp-rate ?u)))
 :condition ( and (at start (running))
       (over all (on ?u))
       (over all (not (rampingDown ?u)))
       (at start (not (rampingUp ?u)))
       (over all (<= (output ?u)
                     (max-output ?u))))
 :effect ( and (at start (rampingUp ?u))
  (increase (output ?u)
                 (* #t (up-ramp-rate ?u)))
        (increase (supply)
                 (* #t (up-ramp-rate ?u)))
        (at end (not (rampingUp ?u)))))
```

Figure 4: The Durative action `ramp-up`

Ramping up / down The ramping actions are used to increment and decrement the current output of a unit. The effect is computed using the continuous effect and considering the rate of change of the output. An `over all` condition prevents the output to exceed the maximum value.

Cooling The minimum off time is enforced by the action `cool-down`. This action is used only if the unit is to be re-switched on, and if so it does not need to run the whole time the unit is off so there is no need for a duration inequality.

```
(:durative-action cool-down
 :parameters (?u - unit)
 :duration (= ?duration (minOffTime ?u))
 :condition ( and (at start (running))
  (at start (off ?u)) (over all (off ?u)))
 :effect ( and (at end (canSwitchOn ?u))))
```

Figure 5: Durative action `cool-down`

UC Planning Problem

The problem file contains the set of units to be managed with their characteristics. The initial state also includes the Timed Initial Fluents describing the demand profile. A fragment of the problem is shown in Figure 6 where, for sake of space, only one unit is described[3].

[3]The interested reader can ask the authors for the full domain and problem descriptions.

```
(define (problem ucp1) (:domain uc)
(:objects u1 ... - unit)
(:init (off u1)
(= (total-cost) 0)  (= (total-supply) 0)
(= (min-output u1) 30) (= (max-output u1) 75)
(= (up-ramp-rate u1) 6)
(= (down-ramp-rate u1) 6)
(= (switch-on-time u1) 4)
(= (switch-off-time u1) 4)
(= (startup-cost u1) 10)
(= (marginal-cost u1) 40)
(at 0 (= (demand) 0.0))
(at 50 (= (demand) 256.0))
...
(at 1440 (= (demand) 230.0)))
(:goal (and (complete))))
```

Figure 6: Fragment of the UC Problem

Planning Challenges

As mentioned above a strength of the planner is its ability to only reason about necessary changes to meet the new demand and leave some characteristics in a steady state as opposed to the MIP reasoning about the binary on / off variables or outputs which remain constant at every time point. On the other hand, one of the challenges is the need of finding an effective heuristic in order to scale and to be able to deal with the huge state space generated by this problem.

Furthermore, the planner must favour plans that are cheap given how far through the planning horizon they have planned for. The current cost of a plan which has only scheduled a few periods is clearly much less than a plan much further through the planning horizon. It is important that the heuristic recognises there is a large cost to go, to intelligently 'guess' what that cost may be so as not to get stuck only advancing plans in the first hour of the planning horizon.

Finally, the heuristic needs to balance speed and quality of the solution. By always selecting the active plan furthest through the planning horizon a poor quality solution can be found very quickly. Imposing a limit on the maximum overhead (supply minus demand) does not ensure a least cost plan as it does not take into account the operating costs of the units in the plan. Some notion of future cost to go, as mentioned above, could help the planner to choose which of the current partial plans to take forward.

Conclusion

Unit Commitment as a MIP is a much studied problem of great importance to the power systems engineering community. AI Planning could potentially advance the work on UC by providing faster solutions, which would be beneficial in light of the upcoming changes to the UC problem. The problem contains a novel collection of constraints which mean all existing planners we are aware of cannot directly support this model. Solving this domain competitively compared to the well established MIP model will advance the field of Planning with time and resource greatly.

Acknowledgments

This work was partially funded by the EPSRC Projects EP/I031650/1 and EP/G023360/2.

References

Agency, I. E. 2012. World Energy Outlook 2009. [Online] : http://www.iea.org/textbase/nppdf/free/2009/weo2009.pdf.

Barth, R.; Brand, H.; Meibom, P.; and Weber, C. 2006. A stochastic unit-commitment model for the evaluation of the impacts of integration of large amounts of intermittent wind power. In *9th International Conference on Probabilistic Methods Applied to Power Systems*. PMAPS.

Bell, K.; Coles, A. I.; Fox, M.; Long, D.; and Smith, A. J. 2008. The application of planning to power substation voltage control. In *ICAPS Workshop on Scheduling and Planning Applications (SPARK)*.

Carrión, M., and Arroyo, J. 2006. A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem. *Proceedings of the IEEE* 21(3):1371–1378.

Coles, A.; Coles, A.; Fox, M.; and Long, D. 2012. COLIN : Planning with continuous linear numeric change. *Journal of Artificial Intelligence Research* 44:1–96.

Commission, E. 2012. [Online] : http://ec.europa.eu/clima/policies/package/index_en.htm.

Consulting, P. E. 2009. Impact of intermittency: How wind variability could change the shape of the British and Irish electricity markets. [Online] : http://www.uwig.org/ImpactofIntermittency.pdf.

Della Penna, G.; Intrigila, B.; Magazzeni, D.; and Mercorio., F. 2009. UPMurphi: a tool for universal planning on PDDL+ problems. In *Proceedings of the 19th International Conference on Automated Planning and Scheduling*. ICAPS.

Fox, M., and Long, D. 2003. An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20:61–124.

Fox, M.; Long, D.; and Magazzeni, D. 2011. Automatic construction of efficient multiple battery usage policies. In *International Conference on Automated Planning and Scheduling (ICAPS)*.

Fox, M.; Long, D.; and Magazzeni, D. 2012. Plan-based policies for efficient multiple battery load management. *J. Artif. Intell. Res. (JAIR)* 44:335–382.

Ilic, M. 2007. From hierarchical to open access electric power systems. *Proceedings of the IEEE* 95(5):1060 –1084.

Meibom, P.; Barth, R.; Brand, H.; and Weber, C. 2007. Wind power integration studies using a multi-stage stochastic electricity system model. In *Power Engineering Society General Meeting*. IEEE.

Meibom, P.; Barth, R.; Hasche, B.; Brand, H.; Weber, C.; and OMalley, M. 2011. Stochastic optimization model to study the operational impacts of high wind penetrations in Ireland. *IEEE Transactions On Power Systems* 26(3):1367–1379.

on Climate Change, T. C. 2012. [Online] : http://www.theccc.org.uk/sectors/power/scenarios.

Ostrowski, J.; Anjos, M.; and Vannelli, A. 2012. Tight mixed integer linear programming formulations for the unit commitment problem. *IEEE Transactions on Power Systems* 27(1):39–46.

RenewableUK. 2012. [Online] : http://www.renewableuk.com/en/renewable-energy/wind-energy/uk-wind-energy-database/index.cfm.

Streiffert, D.; Philbrick, R.; and Ott, A. 2005. A mixed integer programming solution for market clearing and reliability analysis. In *Power Engineering Society General Meeting*. IEEE.

Sturt, A., and Strbac, G. 2012. Efficient stochastic scheduling for simulation of wind-integrated power systems. *IEEE Transactions on Power Systems* 27(1):323–334.

Tuohy, A.; Meibom, P.; Denny, E.; and OMalley, M. 2009. Unit commitment for systems with significant wind penetration. *IEEE Transactions on Power Systems* 24(2):592–601.

Wollenberg, B., and Wood, A. 1996. *Power Generation, Operation and Control*. John Wiley, 2nd edition.