Energy-Cost Aware Scheduling/Forecasting Competition

H. Simonis, B. O'Sullivan, D. Mehta, B. Hurley, M. De Cauwer

June 17, 2014

Abstract

In this note we describe the setting, the rules and the format for a forecasting/scheduling competition about energy cost aware cloud computing, planned for CP-AI-OR 2014 in Cork and CP2014 in Lyon. We discuss the forecasting and the scheduling problem, present a mathematical formulation of the problem, and define the data format(s) used. Participants can decide to participate by providing either a price forecast evaluated on an existing scheduling system, a scheduling system based on a given price forecast, or by providing a combination of forecasting and scheduling tools. The quality of the solutions is evaluated through the total, actual costs of the schedules produced.

1 Setting

We consider the following problem: You are running a cloud computing service, where customers contract to run computing services (tasks). Each task has a duration, an earliest start and latest end, and resource requirements expressed as integer values for CPU, Memory and I/O attributes. The tasks can be scheduled on one of multiple servers, each server has a limited capacity for the CPU, memory and I/O attributes. Multiple tasks can run concurrently on the same machine if the total resource consumption for all attributes is below the respective capacity. All tasks must be scheduled within their release and due dates, these dates are set so that no task stretches over midnight between two days. Tasks can not be interrupted, once started, they must run for their given duration.

If a machine is used by a task, it must be running at that time. In addition to the cost of running the allocated tasks, the machine consumes some idle power if it is on. Every time a machine is switched on or off, a start-up resp. shut-down cost must be paid. All machines are off at the beginning of the planning period, all machines must be off at the end of the planning period.

The price of electricity for the data centre is a real-time price, and varies throughout the day. The actual price is not known in advance, a forecast must be used to generate a schedule. The total cost of the schedule is determined after the fact by applying the actual price of electricity to the energy consumption in each time period. One forecast of the price is given by the organizers. An example of the discrepancy between the forecast and actual price is shown in Figure 1, offering the opportunity to generate better forecasts based on historical data for demand and prices, and previous forecast information. Note that a forecast with a low error is not automatically guaranteed to lead to a schedule with a low overall cost.

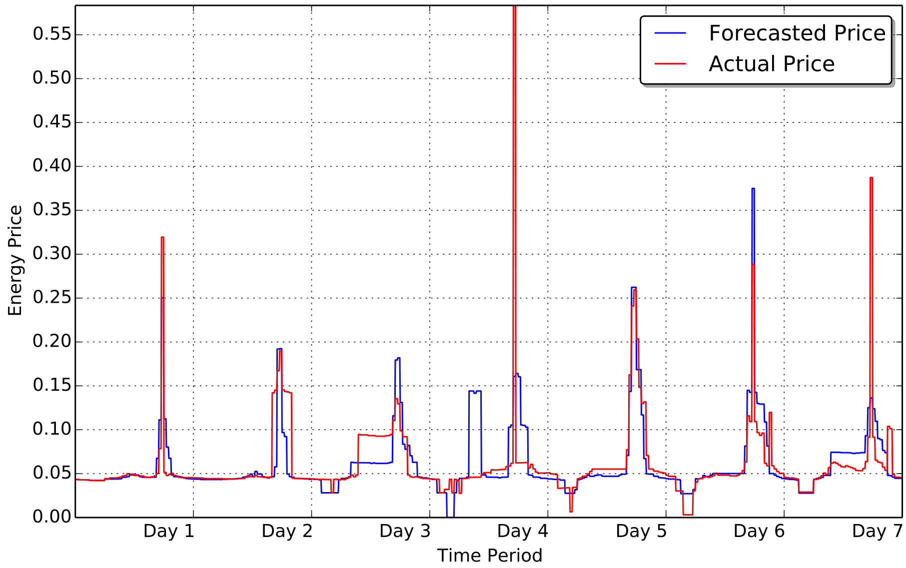

Figure 1: Energy price data from the Irish electricity market for a sample one-week period.

The objective is to generate a schedule with the lowest total actual cost, within a certain time limit.

The competition is run in three categories:

1. You use the provided price forecast to generate your best possible schedule. All tasks must be scheduled, the cost of the schedule is given by the actual price for electricity. The validity and cost of the schedule generated is tested by a solution checker.

2. You generate your own price forecast, and use a provided scheduling system to generate a schedule. The cost of the solution will depend on the quality of your price forecast.

3. You generate both the price forecast and a scheduling system. You provide a schedule, its cost will be determined by a solution checker.

2 Scheduling Problem

In this section we introduce a formal description of the scheduling part of the competition.

2.1 Notation

The following entities, sorted alphabetically, are defined in the model.

c_{mr} resource capacity of machine m for resource r

d_j duration of tasks j

$\mathbf{down}_m$ shut-down cost of server m

e_j earliest start of task j

f_t forecast price of electricity in time period t

$\mathbf{idle}_m$ idle cost of server m

j task index

l_j latest end of task j

m machine index

p_j power use of tasks j

q time resolution

r resource index

r_t actual price of electricity in time period t

t time index

u_{jr} resource use of task j for resource r

$\mathbf{up}_m$ startup cost of server m

v_{mt} server m is running at time t

x_{jmt} task j starts on machine m at time t

y_{mt} server m starts up at time t

z_{mt} server m shuts down at time t

2.2 Sets and Indices

We use the following sets of objects in our model:

J The set of tasks to be scheduled on one day. All tasks given must be scheduled completely within that day, tasks can not be rejected.

T The set of time-periods in a day. We consider a time resolution of q minutes, i.e. there are $1440/q$ time periods in a day.

M the set of servers considered

R the set of resources considered

We also try to use consistent indices in the description:

j index ranging over tasks

t index ranging over time periods

m index ranging over servers (machines)

r index ranging over resources

2.3 Constants

q The time resolution in the scheduling model. The value is an integer expressed in minutes, i.e. there are $|T| = 24 * 60/q + 1$ time periods in a day.

d_j duration of task j, the values are positive integers less than or equal to $|T|$. The duration gives the length of the task, so that the end is the start plus the duration. The task is not considered active at its end-time.

e_j earliest start of task j, the values are integers between 0 and $|T|$.

l_j latest end of task j, the values are integers between 0 and $|T|$. This latest end corresponds to a latest start of $l_j - d_j$.

u_{jr} resource use for task j for resource r, this is a positive integer value.

p_j power use of task j, the power use is constant during execution of the task. This is a non-negative integer value.

c_{mr} capacity of server m for resource r, this is a positive integer.

$\mathbf{idle}_m$ idle power of running server m for one time period, this is a non-negative integer number.

$\mathbf{up}_m$ start-up cost of server m. The cost of starting up machine m once. This cost is not dependent on the energy cost at this time. The value is a non-negative floating point number.

$\mathbf{down}_m$ shut-down cost of server m. The cost of shutting down a server once. This cost is not dependent on the energy cost at the shutdown time. The values is a non-negative floating point number.

f_t This is the forecasted price of one unit of energy in time period t. This is a floating point number, and can be zero or even negative.

r_t the actual price of one unit of energy in time period t. This is a floating point number, and can be zero or even negative. This information can not be used to make scheduling decisions, and is only used to evaluate a resulting schedule.

2.4 Model

We now describe the variables, constraints and objective of the model.

2.4.1 Variables

x_{jmt} task j starts at time t on machine m. A 0/1 integer variable indicating when and where a task is run. As each task must be scheduled, exactly one of the variables linked to a task must be set to one.

v_{mt} server m is running at time t. If any task is running on a machine at some time t, the machine must be active.

y_{mt} machine m starts up at time t. Initially, all machines are off.

z_{mt} machine m shuts down after time period t. The machine is active at time t, but not at time $t+1$. All machines must be switched off at the end of the scheduling horizon.

2.4.2 Constraints

We first need to enforce that each task is scheduled on one machine, exactly once:

$$\forall_{j\in J}: \quad \sum_{m\in M}\sum_{t\in t} x_{jmt} = 1 \tag{1}$$

No task can be scheduled before its earliest start:

$$\forall_{j\in J}\forall_{m\in M}\forall_{t<e_{ij}}: \quad x_{jmt} = 0 \tag{2}$$

No task can be scheduled to end after its latest end:

$$\forall_{j\in J}\forall_{m\in M}\forall_{t+d_{ij}>l_{ij}}: \quad x_{jmt} = 0 \tag{3}$$

The resource requirements of all tasks scheduled on the same machine at the same time must fit within the capacity of the machine:

$$\forall_{m\in M}\forall_{r\in R}\forall_{t\in T}: \quad \sum_{j\in J}\sum_{t-d_j<t'\leq t} x_{jmt'}u_{jr} \leq c_{mr} \tag{4}$$

The following constraints all link the different types of variables related to machines.

If a machine is starting at time t, then we consider it running at this time.

$$\forall_{m\in M}\forall_{t\in T}: \quad y_{mt} \Rightarrow v_{mt} \tag{5}$$

If a machine is starting up at time t, then it was not running at time $t-1$.

$$\forall_{m\in M}\forall_{t\in T}: \quad y_{mt} \Rightarrow v_{mt-1} = 0 \tag{6}$$

If a machine is shutting down at time t, then it is still running at this time.

$$\forall_{m\in M}\forall_{t\in T}: \quad z_{mt} \Rightarrow v_{mt} \tag{7}$$

If a machine is shutting down at time t, then is is not running at time $t+1$.

$$\forall_{m\in M}\forall_{t\in T}: \quad z_{mt} \Rightarrow v_{mt+1} = 0 \tag{8}$$

If a task is starting at time t on some machine m, then that machine must be active at least while the task is running, i.e. for all time points from t to $t+d_j-1$.

$$\forall_{j\in J}\forall_{m\in M}\forall_{t\in T}\forall_{t\leq t'<t+d_j}: \quad x_{jmt} \Rightarrow v_{mt'} \tag{9}$$

If a machine is running at time t, then it was either already running at time $t-1$ or it starts up at time t.

$$\forall_{m\in M}\forall_{t\in T}: \quad v_{mt} \Rightarrow v_{mt-1} \vee y_{mt} \tag{10}$$

If a machine is running at time t, then it is either also running at time $t+1$, or it shuts down at time t.

$$\forall_{m\in M}\forall_{t\in T}: \quad v_{mt} \Rightarrow v_{mt+1} \vee z_{mt} \tag{11}$$

2.4.3 Objective Function

The objective is to minimize the total cost of operation, consisting of the energy cost running all tasks c_J, the cost of running the servers when they are active c_M, the startup cost of the servers c_{up} and the shutdown cost of the servers c_{down}. This means we minimize the following function

$$\text{cost} := \min c_J + c_M + c_{\text{up}} + c_{\text{down}} \tag{12}$$

The energy cost of running all tasks is given by the sum

$$c_J := \sum_{j \in J} \sum_{m \in M} \sum_{t \in T} x_{jmt} \Big(\sum_{t \le t' < t + d_j} p_j r_{t'} q/60 \Big) \tag{13}$$

Note that we have to convert the power use for the task into an energy value by multiplying with the duration of the time period (in hours). As we don't know the actual cost of electricity r_t when creating the schedule, we may decide to use the forecast price f_t in the optimisation instead. This may mean that an optimal solution for the forecast is not optimal for the actual price. But the final evaluation of the solution quality will be based on the actual price, which is only known after the fact.

The energy cost of running the servers (ignoring the cost of the tasks) is given by

$$c_M := \sum_{m \in M} \sum_{t \in T} v_{mt} \text{idle}_m r_t q/60 \tag{14}$$

The start-up cost is given by the sum

$$c_{\text{up}} := \sum_{m \in M} \sum_{t \in T} y_{mt} \text{up}_m \tag{15}$$

The shut-down cost is given by the sum

$$c_{\text{down}} := \sum_{m \in M} \sum_{t \in T} z_{mt} \text{down}_m \tag{16}$$

2.5 Data Format

The scheduling problem is given in the following text file format. Numbers in parentheses give the typical size of the value.

```
Time_Resolution (5)
Nr_of_Resources n (3)
Nr_of_Machines (>20)
    Machine[k 1..Nr_of_Machines] idle up down
    resource_1 resource_2 ... resource_n
Nr_of_Tasks (> 2000)
    Task[j 1..Nr_of_Tasks] duration earliest_start latest_end power
    resource_1 resource_2 ... resource_n
```

Duration, earliest_start and latest_end are given as integers in multiple of q minutes, with a total range from 0 to $1440/q$. All values except costs are integers, numbers are separated by one or more white space characters.

The price data is given by the following file format:

```
Nr of Periods (288)
    Period[1..nr_of_Periods] cost [euro/kWh]
```

Prices are given as floating point numbers. Forecast price solutions must be generated in the same format.

The schedule is generated, listing the machine and start time of each task:

```
Nr_machines
   Machine[m 1..Nr_Machines]
   Nr_events
      on/off[1/0] time_period
Nr_of_Tasks (> 2000)
    Task[j 1..Nr_of_Tasks] machine start
```

The start times are given as time periods since midnight on the day, in multiples of q.

2.6 Alternative Description of Data Format

This section describes the data formats of the various components. The first section describes the format of the scheduling problem, the second describes the format for the energy price, and the last describes the scheduling solution format. Example data is given next to each section and is indented to signify repeated sections. Prices are given as floating point numbers. All times are given as time periods since midnight on the day, in multiples of q.

2.6.1 Scheduling Problem

The scheduling problem is given in the following text file format and illustrated with some sample values. The first two lines contain a single integer each representing q and the number of resources $|R|$ respectively.

```
q                                 5
nr_of_resources                   3
```

The next line of input contains a single integer representing the number of machines $|M|$ available. For each machine, we define two lines of input. The first line contains the identifier, m, followed by three numbers representing idle_m, up_m, and down_m. idle_m is the energy usage rate in Watts of machine m for a time period if it is running. Both up_m and down_m are given as Euro costs. The second line contains $|R|$ integers representing the respective resource capacities c_{mr} of machine m.

```
nr_of_machines                    20
  m idle_m up_m down_m            0 50 1.0 1.0
  c_m0 c_m1 c_m2                  92 71 63
```

The next line contains a single integer $|J|$, the number of tasks to be scheduled. This is followed by two lines for each task. The first of which describes the task ID, d_j, e_j, and l_j. The second line contains the resource requirements for the task u_{jr}.

```
nr_tasks                          100
  task_j d_j e_j l_j p_j          0 13 185 219 142.0
  u_j1 u_j2 u_j3                  16 18 29
```

2.6.2 Energy Price Data

This section describes the format of the energy price forecast that will be provided. Forecast price solutions must be generated in the same format.

The first line contains the number of time periods in the day, which will be the value $1440/q$. The next $1440/q$ lines will contain two numbers: the index of the time periods, followed by a decimal representing the energy price for that time period.

```
nr_of_periods                     288
  t_0 [f_0|r_0]                   0 0.04331
  t_1 [f_1|r_1]                   1 0.04331
```

2.6.3 Scheduling Solution

A solution to the scheduling problem should start with an integer representing the number of machines $|M|$. The next $|M|$ lines should contain the index of the machine and the number of on/off events, `nr_events`, associated with the machine. The next `nr_events` lines should contain two integers, the action and the time period t. The action should be 0 (1) if the machine should be turned off (on) at time period t. All machines are initially off and should be off by the end of the day also.

The next line should contain a single integer representing the number of tasks, $|J|$. The next $|J|$ lines should contain three integers: the index task j, the machine m to which it should be assigned, and the start time of the task on that machine.

```
nr_machines                       20
  machine_m                       0
  nr_events                       2
    on/off[1/0] time_t            1 0
    on/off[1/0] time_t            0 287
nr_of_tasks                       100
  task_j machine start            0 4 256
```

2.6.4 Sample Instance

To illustrate the data format a complete sample instance with 1 machine, 3 resources, and 10 tasks is given below.

```
5
3
1
0 100 1.0 1.0
122 135 107
10
0 8 206 265 0.03
10 9 9
1 5 7 67 0.59
10 9 9
```

```
2 44 83 153 0.10
10 9 9
3 37 202 275 0.55
9 10 9
4 37 216 274 0.29
10 9 9
5 6 28 105 0.21
10 9 9
6 19 46 228 0.64
9 9 10
7 15 230 279 0.27
9 9 9
8 24 173 234 0.25
10 9 9
9 13 160 277 0.51
9 9 10
```

A scheduling solution for this instance is:

```
1
0
2
1 0
0 287
10
0 0 256
1 0 61
2 0 83
3 0 237
4 0 236
5 0 41
6 0 53
7 0 263
8 0 174
9 0 263
```

3 Forecast Problem

In the forecast problem, we have to predict the actual electricity price for one day into the future based on historical and forecasted data. The historical data is available from September 2011 onwards. Missing values are marked with NaN. The following fields are defined:

DateTime String, defines date and time of sample

Holiday String, gives name of holiday if day is a bank holiday

HolidayFlag integer, 1 if day is a bank holiday, zero otherwise

DayOfWeek integer (0-6), 0 monday, day of week

WeekOfYear integer, running week within year of this date

Day integer, day of the date

Month integer, month of the date

Year integer, year of the date

PeriodOfDay integer, denotes half hour period of day (0-47)

WindForecast the forecasted wind production for this period

LoadForecast th national load forecast for this period

PriceForecast the price forecast for this period

Temperature the actual temperature measured at Cork airport

Windspeed the actual windspeed measured at Cork airport

CO2Intensity the actual CO2 intensity in (g/kWh) for the electricity produced

ActualWind the actual wind energy production for this period

ActualLoad the actual national system load for this period

ActualPrice the actual price of this time period, the value to be forecasted

The last four fields are only available for historical data, i.e. they can not be used to make the forecast.

The following shows a sample day of the historical data available, a header line starting with # to describe the fields is the first line of the data file:

```
#DateTime Holiday HolidayFlag DayOfWeek WeekOfYear Day Month Year
PeriodOfDay WindForecast LoadForecast PriceForecast Temperature
Windspeed CO2Intensity ActualWind ActualLoad ActualPrice
"mon 31/3/2014 00:00" "" 0 0 14 31 3 2014 0  523.70 3406.49   31.60    8.00   13.00  513.64  516.00 2758.73   28.57
"mon 31/3/2014 00:30" "" 0 0 14 31 3 2014 1  472.80 3283.17   31.60    8.00   11.10  524.15  519.00 3038.27   31.60
"mon 31/3/2014 01:00" "" 0 0 14 31 3 2014 2  430.70 3213.60   31.60    8.00   11.10  532.90  526.00 2926.23   28.60
"mon 31/3/2014 01:30" "" 0 0 14 31 3 2014 3  415.60 3078.08   31.60    8.00   13.00  536.32  489.00 2842.25   31.22
"mon 31/3/2014 02:00" "" 0 0 14 31 3 2014 4  403.40 2999.27   31.60    8.00   11.10  538.74  501.00 2806.94   28.57
"mon 31/3/2014 02:30" "" 0 0 14 31 3 2014 5  401.80 2923.80   31.60    8.00   13.00  537.57  514.00 2744.80   28.63
"mon 31/3/2014 03:00" "" 0 0 14 31 3 2014 6  405.60 2922.17   31.60    8.00   11.10  522.95  528.00 2688.60   28.60
"mon 31/3/2014 03:30" "" 0 0 14 31 3 2014 7  426.70 2889.29   31.60    8.00    9.30  514.68  553.00 2658.32   28.57
"mon 31/3/2014 04:00" "" 0 0 14 31 3 2014 8  447.80 2908.93   31.60    8.00   13.00  503.80  545.00 2649.23   28.60
"mon 31/3/2014 04:30" "" 0 0 14 31 3 2014 9  467.50 2877.61   28.60    8.00   14.80  511.01  534.00 2660.41   28.60
"mon 31/3/2014 05:00" "" 0 0 14 31 3 2014 10  484.40 2913.92   28.60    8.00   16.70  535.02  536.00 2735.55   28.60
"mon 31/3/2014 05:30" "" 0 0 14 31 3 2014 11  491.60 2969.77   31.60    8.00   16.70  553.32  590.00 2844.59   31.22
"mon 31/3/2014 06:00" "" 0 0 14 31 3 2014 12  500.10 3125.12   33.99    8.00   16.70  597.89  609.00 3022.15   39.37
"mon 31/3/2014 06:30" "" 0 0 14 31 3 2014 13  508.30 3305.45   37.57    8.00   16.70  613.07  619.00 3273.79   39.40
"mon 31/3/2014 07:00" "" 0 0 14 31 3 2014 14  521.50 3634.30   39.37    8.00   18.50  610.97  618.00 3653.93   39.77
"mon 31/3/2014 07:30" "" 0 0 14 31 3 2014 15  543.10 4005.15   56.24    8.00   16.70  587.80  597.00 4021.03   45.91
"mon 31/3/2014 08:00" "" 0 0 14 31 3 2014 16  562.10 4251.82   56.24    8.00   18.50  543.71  611.00 4305.86   46.18
"mon 31/3/2014 08:30" "" 0 0 14 31 3 2014 17  571.90 4371.62   45.40    8.00   22.20  507.77  694.00 4394.86   46.00
"mon 31/3/2014 09:00" "" 0 0 14 31 3 2014 18  578.70 4502.86   45.40    8.00   24.10  503.42  730.00 4457.45   42.25
"mon 31/3/2014 09:30" "" 0 0 14 31 3 2014 19  576.80 4641.48   45.40    8.00   20.40  495.43  848.00 4516.30   42.25
"mon 31/3/2014 10:00" "" 0 0 14 31 3 2014 20  574.20 4732.00   45.40    8.00   20.40  499.29  809.00 4492.21   42.25
"mon 31/3/2014 10:30" "" 0 0 14 31 3 2014 21  569.90 4758.16   45.40    8.00   22.20  504.62  812.00 4494.85   45.82
"mon 31/3/2014 11:00" "" 0 0 14 31 3 2014 22  564.00 4801.00   45.40    8.00   20.40  501.14  759.00 4532.89   45.94
"mon 31/3/2014 11:30" "" 0 0 14 31 3 2014 23  553.20 4823.92   45.40    9.00   18.50  495.47  770.00 4530.71   45.93
"mon 31/3/2014 12:00" "" 0 0 14 31 3 2014 24  543.10 4816.22   62.77    9.00   18.50  494.27  791.00 4552.54   46.05
"mon 31/3/2014 12:30" "" 0 0 14 31 3 2014 25  530.30 4835.71   62.77    9.00   24.10  491.04  753.00 4562.43   46.19
"mon 31/3/2014 13:00" "" 0 0 14 31 3 2014 26  518.00 4824.78   62.77    9.00   20.40  499.68  714.00 4488.21   46.04
"mon 31/3/2014 13:30" "" 0 0 14 31 3 2014 27  505.00 4710.76   62.77   10.00   22.20  494.48  711.00 4402.01   45.76
"mon 31/3/2014 14:00" "" 0 0 14 31 3 2014 28  493.00 4693.38   62.77   10.00   24.10  501.18  667.00 4404.51   45.95
"mon 31/3/2014 14:30" "" 0 0 14 31 3 2014 29  484.00 4681.50   62.77   10.00   22.20  501.45  662.00 4387.11   45.89
"mon 31/3/2014 15:00" "" 0 0 14 31 3 2014 30  474.40 4679.47   62.77   11.00   18.50  501.24  651.00 4402.30   46.06
"mon 31/3/2014 15:30" "" 0 0 14 31 3 2014 31  463.30 4716.79   62.77   10.00   20.40  503.18  587.00 4451.41   46.19
"mon 31/3/2014 16:00" "" 0 0 14 31 3 2014 32  453.20 4804.95   62.77   11.00   20.40  511.02  560.00 4536.09   60.44
"mon 31/3/2014 16:30" "" 0 0 14 31 3 2014 33  441.90 4913.97   62.77   11.00   22.20  505.86  502.00 4643.65   91.33
"mon 31/3/2014 17:00" "" 0 0 14 31 3 2014 34  432.40 5021.85   62.77   10.00   18.50  525.67  454.00 4723.89   91.37
"mon 31/3/2014 17:30" "" 0 0 14 31 3 2014 35  420.30 5094.42   82.38   10.00   16.70  516.07  406.00 4736.14   91.37
"mon 31/3/2014 18:00" "" 0 0 14 31 3 2014 36  406.10 5120.89   82.85   10.00   18.50  521.39  425.00 4635.03   91.37
"mon 31/3/2014 18:30" "" 0 0 14 31 3 2014 37  384.90 5224.30   82.85    9.00   16.70  531.57  392.00 4551.48   91.37
"mon 31/3/2014 19:00" "" 0 0 14 31 3 2014 38  500.60 5376.49  290.46    9.00   14.80  532.84  312.00 4517.69   91.37
"mon 31/3/2014 19:30" "" 0 0 14 31 3 2014 39  450.30 5360.82   82.85    9.00   13.00  548.86  291.00 4565.25   91.47
"mon 31/3/2014 20:00" "" 0 0 14 31 3 2014 40  399.00 5245.61   82.85    8.00   11.10  536.88  267.00 4745.85   91.84
"mon 31/3/2014 20:30" "" 0 0 14 31 3 2014 41  344.20 5061.93   62.77    8.00   11.10  536.09  255.00 4845.91   91.84
"mon 31/3/2014 21:00" "" 0 0 14 31 3 2014 42  296.00 4933.72   45.40    8.00   13.00  543.71  264.00 4730.18   91.84
"mon 31/3/2014 21:30" "" 0 0 14 31 3 2014 43  265.70 4697.79   45.40    8.00   11.10  566.28  254.00 4494.18   91.84
"mon 31/3/2014 22:00" "" 0 0 14 31 3 2014 44  240.00 4411.35   41.88    8.00   13.00  578.28  236.00 4249.64   60.45
```

```
"mon 31/3/2014 22:30" "" 0 0 14 31 3 2014 45  229.70 4150.57   41.88    8.00   13.00  588.65  250.00 3958.54   60.04
"mon 31/3/2014 23:00" "" 0 0 14 31 3 2014 46  220.80 3940.90   41.88    8.00   11.10  609.46  228.00 3694.28   56.51
"mon 31/3/2014 23:30" "" 0 0 14 31 3 2014 47  217.80 3885.61   41.88    8.00   11.10  628.78  201.00 3455.13   54.03
```

For the forecast problem, only the following data will be given:

```
#DateTime Holiday HolidayFlag DayOfWeek WeekOfYear Day Month Year
PeriodOfDay WindForecast LoadForecast PriceForecast Temperature
Windspeed
"mon 31/3/2014 00:00" "" 0 0 14 31 3 2014 0  523.70 3406.49   31.60    8.00   13.00
"mon 31/3/2014 00:30" "" 0 0 14 31 3 2014 1  472.80 3283.17   31.60    8.00   11.10
"mon 31/3/2014 01:00" "" 0 0 14 31 3 2014 2  430.70 3213.60   31.60    8.00   11.10
"mon 31/3/2014 01:30" "" 0 0 14 31 3 2014 3  415.60 3078.08   31.60    8.00   13.00
"mon 31/3/2014 02:00" "" 0 0 14 31 3 2014 4  403.40 2999.27   31.60    8.00   11.10
"mon 31/3/2014 02:30" "" 0 0 14 31 3 2014 5  401.80 2923.80   31.60    8.00   13.00
"mon 31/3/2014 03:00" "" 0 0 14 31 3 2014 6  405.60 2922.17   31.60    8.00   11.10
"mon 31/3/2014 03:30" "" 0 0 14 31 3 2014 7  426.70 2889.29   31.60    8.00    9.30
"mon 31/3/2014 04:00" "" 0 0 14 31 3 2014 8  447.80 2908.93   31.60    8.00   13.00
"mon 31/3/2014 04:30" "" 0 0 14 31 3 2014 9  467.50 2877.61   28.60    8.00   14.80
"mon 31/3/2014 05:00" "" 0 0 14 31 3 2014 10  484.40 2913.92   28.60    8.00   16.70
"mon 31/3/2014 05:30" "" 0 0 14 31 3 2014 11  491.60 2969.77   31.60    8.00   16.70
"mon 31/3/2014 06:00" "" 0 0 14 31 3 2014 12  500.10 3125.12   33.99    8.00   16.70
"mon 31/3/2014 06:30" "" 0 0 14 31 3 2014 13  508.30 3305.45   37.57    8.00   16.70
"mon 31/3/2014 07:00" "" 0 0 14 31 3 2014 14  521.50 3634.30   39.37    8.00   18.50
"mon 31/3/2014 07:30" "" 0 0 14 31 3 2014 15  543.10 4005.15   56.24    8.00   16.70
"mon 31/3/2014 08:00" "" 0 0 14 31 3 2014 16  562.10 4251.82   56.24    8.00   18.50
"mon 31/3/2014 08:30" "" 0 0 14 31 3 2014 17  571.90 4371.62   45.40    8.00   22.20
"mon 31/3/2014 09:00" "" 0 0 14 31 3 2014 18  578.70 4502.86   45.40    8.00   24.10
"mon 31/3/2014 09:30" "" 0 0 14 31 3 2014 19  576.80 4641.48   45.40    8.00   20.40
"mon 31/3/2014 10:00" "" 0 0 14 31 3 2014 20  574.20 4732.00   45.40    8.00   20.40
"mon 31/3/2014 10:30" "" 0 0 14 31 3 2014 21  569.90 4758.16   45.40    8.00   22.20
"mon 31/3/2014 11:00" "" 0 0 14 31 3 2014 22  564.00 4801.00   45.40    8.00   20.40
"mon 31/3/2014 11:30" "" 0 0 14 31 3 2014 23  553.20 4823.92   45.40    9.00   18.50
"mon 31/3/2014 12:00" "" 0 0 14 31 3 2014 24  543.10 4816.22   62.77    9.00   18.50
"mon 31/3/2014 12:30" "" 0 0 14 31 3 2014 25  530.30 4835.71   62.77    9.00   24.10
"mon 31/3/2014 13:00" "" 0 0 14 31 3 2014 26  518.00 4824.78   62.77    9.00   20.40
"mon 31/3/2014 13:30" "" 0 0 14 31 3 2014 27  505.00 4710.76   62.77   10.00   22.20
"mon 31/3/2014 14:00" "" 0 0 14 31 3 2014 28  493.00 4693.38   62.77   10.00   24.10
"mon 31/3/2014 14:30" "" 0 0 14 31 3 2014 29  484.00 4681.50   62.77   10.00   22.20
"mon 31/3/2014 15:00" "" 0 0 14 31 3 2014 30  474.40 4679.47   62.77   11.00   18.50
"mon 31/3/2014 15:30" "" 0 0 14 31 3 2014 31  463.30 4716.79   62.77   10.00   20.40
"mon 31/3/2014 16:00" "" 0 0 14 31 3 2014 32  453.20 4804.95   62.77   11.00   20.40
"mon 31/3/2014 16:30" "" 0 0 14 31 3 2014 33  441.90 4913.97   62.77   11.00   22.20
"mon 31/3/2014 17:00" "" 0 0 14 31 3 2014 34  432.40 5021.85   62.77   10.00   18.50
"mon 31/3/2014 17:30" "" 0 0 14 31 3 2014 35  420.30 5094.42   82.38   10.00   16.70
"mon 31/3/2014 18:00" "" 0 0 14 31 3 2014 36  406.10 5120.89   82.85   10.00   18.50
"mon 31/3/2014 18:30" "" 0 0 14 31 3 2014 37  384.90 5224.30   82.85    9.00   16.70
"mon 31/3/2014 19:00" "" 0 0 14 31 3 2014 38  500.60 5376.49  290.46    9.00   14.80
"mon 31/3/2014 19:30" "" 0 0 14 31 3 2014 39  450.30 5360.82   82.85    9.00   13.00
"mon 31/3/2014 20:00" "" 0 0 14 31 3 2014 40  399.00 5245.61   82.85    8.00   11.10
"mon 31/3/2014 20:30" "" 0 0 14 31 3 2014 41  344.20 5061.93   62.77    8.00   11.10
"mon 31/3/2014 21:00" "" 0 0 14 31 3 2014 42  296.00 4933.72   45.40    8.00   13.00
"mon 31/3/2014 21:30" "" 0 0 14 31 3 2014 43  265.70 4697.79   45.40    8.00   11.10
"mon 31/3/2014 22:00" "" 0 0 14 31 3 2014 44  240.00 4411.35   41.88    8.00   13.00
"mon 31/3/2014 22:30" "" 0 0 14 31 3 2014 45  229.70 4150.57   41.88    8.00   13.00
"mon 31/3/2014 23:00" "" 0 0 14 31 3 2014 46  220.80 3940.90   41.88    8.00   11.10
"mon 31/3/2014 23:30" "" 0 0 14 31 3 2014 47  217.80 3885.61   41.88    8.00   11.10
```

As a solution, the predicted price should be given:

```
#DateTime ActualPrice
"mon 31/3/2014 00:00" 28.57
"mon 31/3/2014 00:30" 31.60
"mon 31/3/2014 01:00" 28.60
"mon 31/3/2014 01:30" 31.22
"mon 31/3/2014 02:00" 28.57
"mon 31/3/2014 02:30" 28.63
"mon 31/3/2014 03:00" 28.60
"mon 31/3/2014 03:30" 28.57
"mon 31/3/2014 04:00" 28.60
"mon 31/3/2014 04:30" 28.60
"mon 31/3/2014 05:00" 28.60
"mon 31/3/2014 05:30" 31.22
"mon 31/3/2014 06:00" 39.37
```

```
"mon 31/3/2014 06:30" 39.40
"mon 31/3/2014 07:00" 39.77
"mon 31/3/2014 07:30" 45.91
"mon 31/3/2014 08:00" 46.18
"mon 31/3/2014 08:30" 46.00
"mon 31/3/2014 09:00" 42.25
"mon 31/3/2014 09:30" 42.25
"mon 31/3/2014 10:00" 42.25
"mon 31/3/2014 10:30" 45.82
"mon 31/3/2014 11:00" 45.94
"mon 31/3/2014 11:30" 45.93
"mon 31/3/2014 12:00" 46.05
"mon 31/3/2014 12:30" 46.19
"mon 31/3/2014 13:00" 46.04
"mon 31/3/2014 13:30" 45.76
"mon 31/3/2014 14:00" 45.95
"mon 31/3/2014 14:30" 45.89
"mon 31/3/2014 15:00" 46.06
"mon 31/3/2014 15:30" 46.19
"mon 31/3/2014 16:00" 60.44
"mon 31/3/2014 16:30" 91.33
"mon 31/3/2014 17:00" 91.37
"mon 31/3/2014 17:30" 91.37
"mon 31/3/2014 18:00" 91.37
"mon 31/3/2014 18:30" 91.37
"mon 31/3/2014 19:00" 91.37
"mon 31/3/2014 19:30" 91.47
"mon 31/3/2014 20:00" 91.84
"mon 31/3/2014 20:30" 91.84
"mon 31/3/2014 21:00" 91.84
"mon 31/3/2014 21:30" 91.84
"mon 31/3/2014 22:00" 60.45
"mon 31/3/2014 22:30" 60.04
"mon 31/3/2014 23:00" 56.51
"mon 31/3/2014 23:30" 54.03
```

Note that the forecasting problem uses a different time resolution than the scheduling problem, an automated interface between the two formats will be provided.

4 Evaluation

Registration and the public leaderboard are available on the challenge site: `http://iconchallenge.insight-centre.org`

Each team is limited to one solver submission per category but you may submit solutions to the online leaderboard up to three times per instance per day. Individuals may not participate in multiple teams. Members of the ICON group may participate but their teams are not eligible for prizes.

Leaderboard A leaderboard will be published online throughout the competition. For this, teams need only to submit the solution output for any instances they have solved. Teams are encouraged to regularly submit their solutions throughout the competition. A prize will be awarded to the team which is ranked 1st for the longest between 9th June and 15th August 2014.

Dates Submissions to the public-leaderboard will open on the 9th June 2014; the closing date for leaderboard submissions is the 15th August 2014 23:59 UTC-12. Teams will be required to submit their source code by 17th August 2014 UTC-12, which will be evaluated on 50 unseen instances. These instances will be similar in nature to the public-leaderboard set. The presentation of results will take place at CP 2014, Lyon, France, 8th – 12th September.

Environment Submissions will be run on machines with two Intel Xeon Processor E5-2640 2.5GHz (12 cores, 24 threads) and 64GB RAM. If the solution requires any third party tools, they must be freely available, at least under an academic license. Solutions will be limited to 5 minutes (300 seconds) of wall-clock time per instance for the schedule only and forecast only categories, for the combined forecast and scheduling category this limit is 10 minutes (600 seconds) wall-clock time. These limits only apply to the offline evaluation and not to the public-leaderboard submissions.

Ranking Teams are ranked based on the actual energy cost across all instances. Entries to the forecast-only category will be used to feed a reference scheduler and will be ranked by the schedule produced. Solver run-time is not considered in the rankings, most important is to solve an instance within the time-limit, secondary to this is the relative quality of the solution. Lower is better in this scoring system.

The objective function is computed for each instance using the actual energy price data, not the forecasted energy price. If no solver provides a solution to an instance, that instance is not considered for scoring. If a solver fails to solve an instance but other solvers do, it is given a score equal to the number of solvers that did solve the instance. This penalises solvers relative to the difficulty of the instance. If only one solver solves an instance, it is given a score of zero.

In the case where multiple solvers solve an instance, their relative scores are computed as follows. Let S be the set of solutions for the instance, where $s_i \in S$ is the objective value (using the actual energy price data) of solver i. If all objective values are equivalent, then each solver scores zero, otherwise let $\alpha = \min_{s_j \in S} s_j$ be the best, and $\beta = \max_{s_j \in S} s_j$ be the worst objective value across the solvers. Then, the relative score achieved by solver i for the instance is given as:

$$\text{score}_i = \frac{s_i - \alpha}{\beta - \alpha}$$

Table 1 illustrates this scoring mechanism with an example. In this case, solver A is best with a score of 1.0, followed by B with 2.02, and C with 2.50.

Solver	Instances				Total Score
A	66 (0.00)	100 (0.00)	2000 (0.00)	- (1.00)	(1.00)
B	100 (1.00)	101 (0.02)	3000 (1.00)	1000 (0.00)	(2.02)
C	66 (0.00)	150 (1.00)	2500 (0.50)	- (1.00)	(2.50)

Table 1: An example of the scores achieved by solvers A, B, and C across 4 instances. The values represent their respective objective values for that instance, bracketed values represent the points awarded, and a dash signifies that no solution was given by the solver.