

# Putting everything together - a staff scheduling problem

**Learning objective**

Use OPL scheduling keywords and syntax to model a simple staff scheduling problem

**Key terms**

- alternative resource
- resource pool
- surrogate constraint

**The business problem**

You are now going to use OPL to create a model representing a staff scheduling problem. This is a classic type of task-based scheduling problem to model.

A telephone company must schedule customer requests for installation of different types of telephone lines:

- First (or principal) line
- Second (or additional) line
- ISDN (digital) line

Each request has a requested due date; a due date can be missed, but the objective is to minimize the number of days late.

These three request types each have a list of tasks that must be completed in order to complete the request. There are precedence constraints associated with some of the tasks. Each task has a fixed duration and also may require certain fixed quantities of specific types of resources.

The resource types are

- Operator
- Technician
- CherryPicker (a type of crane)
- ISDNPacketMonitor
- ISDNTechnician.

The tasks types, along with their durations and resource requirements are outlined in the following table:

Task type	Duration	Resources
MakeAppointment	1	Operator
FlipSwitch	1	Technician
InteriorSiteCall	3	Technician x 2
ISDNInteriorSiteCall	3	<ul style="list-style-type: none"> <li>• Technician</li> <li>• ISDNTechnician</li> <li>• ISDNPacketMonitor</li> </ul>
ExteriorSiteCall	2	<ul style="list-style-type: none"> <li>• Technician x 2</li> <li>• CherryPicker</li> </ul>
TestLine	1	Technician

Each request type has a set of task types that must be executed. Some of the tasks must be executed before other tasks can start:

Request type	Task type	Preceding tasks
FirstLineInstall	FlipSwitch	
	TestLine	FlipSwitch

Request type	Task type	Preceding tasks
SecondLineInstall	MakeAppointment	
	InteriorSiteCall	MakeAppointment
	TestLine	InteriorSiteCall
ISDNInstall	MakeAppointment	
	ISDNInteriorSiteCall ExteriorSiteCall	MakeAppointment
	TestLine	InteriorSiteCall, ExteriorSiteCall

### The data

The following resources are available:

Resource	ID/Name
Operator	"Patrick"
Technician	"JohnTec"
Technician	"PierreT"
CherryPicker	"VIN43CP"
CherryPicker	"VIN44CP"
ISDNPacketMonitor	"EQ12ISD"
ISDNTechnician	"RogerTe"

The requests that need to be scheduled are:

Request Number	Request Type	Due date
0	FirstLineInstall	22
1	SecondLineInstall	22
2	FirstLineInstall	1
3	ISDNInstall	21
4	SecondLineInstall	21
5	ISDNInstall	22

# Model the staff scheduling problem



You are now going to perform a series of steps in the **Staff Scheduling** workshop based on this problem.

## What is the objective?

In business terms, we might say the objective is “to improve on-time performance.” However, this is a qualitative statement that is difficult to model. In order to find a modeling representation (i.e. quantifiable) of the idea, we need to turn things around; instead of maximizing something abstract, we find a number that needs to be kept to a minimum. Thus, our objective becomes:

To minimize the total number of late days (days beyond the due date when requests are actually finished).

## What are the unknowns?

The unknowns are:

- When each task will start
- Which resource will be assigned to each task

Before you begin to work with the lab files, here is an overview of what you are going to do. Each task will be decomposed step by step to demonstrate the process.

## How to model this situation:

1. The task type **FlipSwitch** requires a Technician, and there are two Technicians. For each task of this type, you create three **interval** decision variables. Two of these intervals are optional, meaning they may or may not appear in the solution.
2. To constrain such that when an optional interval is present, it is scheduled at exactly the same time as the task interval, use a **synchronize** constraint.
3. To ensure that the appropriate number of worker intervals are used, write a constraint that requires that the sum of that the sum of present worker intervals is equal to the number of resources required.

## Details of the intervals for FlipSwitch

- One optional represents **JohnTec** being assigned to the task
- The other optional represents **Pierret** being assigned to the task.
- The third interval represents the task itself, and is used in other constraints.



In general, if there are multiple resources with identical properties, it is best to model them as a **resource pool**. However, one can imagine that in this example new constraints related to workers' days-off could be added, so here each worker is treated as an individual resource.

Resource pools are explained in the section on **surrogate constraints**.

## Modeling the precedence constraints

To model that some tasks in a request must occur before other tasks in the same request, you use the precedence constraint **endBeforeStart**.

While the data for this problem does not require there to be any delay between tasks, you can add a delay to the model to allow for the possibility of a delay.

In the **Staff Scheduling** workshop, perform the **Declare task interval and precedences** step.

Now you will begin the hands-on part of this exercise, and build the staff scheduling model step by step.

You can perform this lab using the HTML workshop, or by following the instructions in the workbook. The HTML workshop will give you direct access to OPL documentation pages that can help you with the lab.

## Declare task interval and precedences

### Objective

- Start building a scheduling model using some basic CP Optimizer constructs.

### Actions

- Examine data and model files
- Define the tasks and precedences
- Solve the model and examine the output

### References

*interval*  
*endBeforeStart*

### Examine data and model files

1. Import the `sched_staffWork` project into the OPL Projects navigator (Leave the **Copy projects into workspace** box unchecked) and open the `step1.mod` and `data.dat` files.

The `.mod` file represents a part of what the finished model will look like. Most of the model is already done.



Note that there is no objective function. At this point, you have what is called a **satisfiability problem**. Running it will determine values that satisfy the constraints, without the solution necessarily being optimal.

2. Examine closely how the data declarations for the model are formulated.
3. Note especially, the declaration of the set `demands`. This creates a set whose members come from the tuple `Demand`, which is a tuple of tuples (`RequestDat` and `TaskDat`). This set is made sparse by filtering it such that only task/request pairs that are found in the tuple set `recipes` are included. Effectively, it creates a sparse set of required tasks to be performed in a request and operations (the same tasks associated with a given resource). Only valid combinations are in the set.



This is a good example of the power of tuple sets to create sparse sets of complex data.

4. The file `sched_staff.dat` instantiates the data as outlined in the problem definition. it instantiates
  - `ResourceTypes`
  - `RequestTypes`
  - `TaskTypes`
  - `resources`
  - `requests`
  - `tasks`
  - `recipes`
  - `dependencies`
  - `requirements`

Discuss how model and data files are related with your instructor and fellow students.

### Define the tasks and precedences

You are now ready to start declaring decision variables and constraints. At this point, we will define only the tasks, and the precedence rules that control them.

1. Declare an interval decision value to represent the time required to do each request/operation pair in the set **demands**. Name the decision variable **titasks**:

```
dvar interval titasks[d in demands] size d.task.ptime;
```

2. An important aspect of the modeling is expressing the precedence constraints on the tasks (demands). These constraints can be expressed using the constraint **endBeforeStart**.

The **step1.mod** file already contains the preparatory declarations:

```
forall(d1, d2 in demands, dep in dependencies :  

  d1.request == d2.request &&  

  dep.taskb == d1.task.type &&  

  dep.taska == d2.task.type)
```

Examine these declarations with your instructor to understand clearly what they mean.

3. Write the **endBeforeStart** constraint.
4. Compare with the solution in  
`<trainingDir>\OPL63.labs\Scheduling\Staff\solution\sched_staffSolution\step1.mod`

### Solve the model and examine the output

1. Solve the model by right clicking the **Step1** run configuration and selecting **Run this** from the context menu.
2. Look at the results in the **Solutions** output tab. Can you determine what the displayed values represent?
3. Look at the **Engine log** and **Statistics** output tabs, and note that this model is, for the moment, noted as a “Satisfiability problem.”
4. Close **Step1.mod**.

### Modeling the objective

To model the objective you need to determine the end time of each request. The request itself can be seen as an interval with a variable length. The request interval must cover, or span, all the intervals associated with the tasks that comprise the request.

You create the objective by finding the difference between the end time and the due date and minimizing it.

Now perform the next step of the **Staff Scheduling** workshop, **Compute the end of a task and define the objective** and the first substep of the **Define the resource constraints** step: **Review the needs**.

## Compute the end of a task and define the objective

### Actions

- *Compute the time needed for each request*
- *Transform the business objective into the objective function*
- *Solve the model and examine the output*

### References

*span*  
*all*  
*maxl*

## Compute the time needed for each request

1. Open `Step2.mod` for editing.
2. The requests are modeled as interval decision variables. Write the following declaration in the model:

```
dvar interval tirequests[requests];
```

3. Write a `span` constraint to link this decision variable to the appropriate `titasks` instances.



Use the `all` quantifier to associate the required tasks for each request with the appropriate duration.

4. Check your solution against  
`<trainingDir>\OPL63.labs\Scheduling\Staff\solution\sched_staffSolution\Step2.mod`

## Transform the business objective into the objective function

The business objective requires the model to minimize the total number of late days (days beyond the due date when requests are actually finished). To do this in the model, you need to write an objective function that minimizes the time for each request that exceeds the due date.

1. Calculate the number of late days for each request:
  - The data element `requests` is the set of data that instantiates the tuple `RequestDat`. The `duedate` is included in this information.
  - The interval `tirequests` represents the time needed to perform each request.
  - Subtract the `duedate` from the date on which `tirequests` ends.



Use the `endof` function to determine the end time of `tirequests`.

2. Include a test that discards any negative results (requests that finish early) from the objective function.



Use the `max1` function to select the greater of:

- the difference between due date and finish date
- 0

3. Minimize the sum of all the non-negative subtractions, as calculated for each request.

Check the solution in

`<trainingDir>\OPL63.labs\Scheduling\Staff\solution\sched_staffSolution\Step2.mod`

## Solve the model and examine the output

1. Solve the model by right clicking the `Step2` run configuration and selecting **Run this** from the context menu.
2. Look at the **Engine log** and **Statistics** output tabs, and note that the model is now reported as a “Minimization problem,” after the addition of the objective function. Scroll down a little further and notice the number of fails reported.
3. Look at the results in the **Solutions** output tab. You will notice that an objective is now reported, in addition to the values of `titasks` and `tirequests`.

## Review the needs

So far, you have defined the following constraints as identified in the business problem, as outlined in the workbook:

- For each demand task, there are exactly the required number of task-resource **operation** intervals present.

- Each task precedence is enforced.
- Each **request** interval decision variable spans the associated **demand** interval decision variables.

You now need to meet the following needs, not yet dealt with in the model:

- Each task-resource **operation** interval that is present is synchronized with the associated task's **demand** interval.
- There is no overlap in time amongst the present **operation** intervals associated with a given resource.
- At any given time, the number of overlapping task-resource **operation** intervals for a specific resource type do not exceed the number of available resources for that type.

## Modeling the alternative resources

The task type **FlipSwitch** requires a technician. There are two technicians, i.e. there are two **alternative resources**, available to do the same task. You want to be able to optimize how each of these is used relative to the objective.

To create the alternative resources, you declare optional intervals for each possible task/resource pair.

To constrain a resource so that it cannot be used by more than one task at a given time, you need to ensure that there is no overlap in time amongst the intervals associated with a given resource.

### How to model this situation:

1. Create a **sequence** decision variable from those intervals
2. Place a **noOverlap** constraint on the **sequence** decision variable.

Now perform the rest of the **Define the resource constraints** step.

## Assign workers to tasks

It is now time to deal with the question of who does what. We know from the data that there is more than one resource, in some cases, capable of doing a given task. How do we decide who is the best one to send on a particular job?

The key idea in representing a scheduling problem with alternative resources is:

- Model each possible task-resource combination with an optional interval decision variable.
  - Link these with an interval decision variable that represents the entire task itself (using a **synchronize** constraint).
1. Open **Step3.mod** for editing.
  2. You will see that a new data declaration has been added:

```
tuple Operation {
    Demand      dmd;
    ResourceDat resource;
};
{Operation} opers = {<d, r >| d in demands, m in requirements, r
in resources : d.task.type == m.task && r.type == m.resource};
```

The members of the tuple set **opers** are the set of tasks assigned to a resource.

3. There is also a new decision variable associated with this tuple set that calculates the time required for each **operation**:

```
dvar interval tiopers[opers] optional;
```

Note that this variable is optional. If one of the optional interval variables is present in a solution, this indicates that the resource associated with it is assigned to the associated task.



Remember that in this model a task is called a **demand**, and a task-resource pair is called an **operation**.

4. Declare a **sequence** decision variable named **workers**, associated with each resource.



Use **all** to connect each **resource** used in an **operation** to its related **tiopers** duration:

```
dvar sequence workers[r in resources] in all(o in opers :
o.resource == r) tiopers[o];
```

5. Constrain this decision variable using a **noOverlap** constraint to indicate the order in which a resource performs its **operations**.

### “Just enough” constraint

Another constraint states that for each demand task, there are exactly the required number of task-resource operation intervals present (“just enough” to do the job – not more or less). The presence of an optional interval can be determined using the **presenceOf** constraint:

```
forall(d in demands, rc in requirements : rc.task == d.task.type) {
  sum (o in opers : o.dmd == d && o.resource.type == rc.resource)
  presenceOf(tioopers[o]) == rc.quantity;
```

- Write this into the model file.

### Check your work and solve the model

1. Compare your results with the contents of  
<trainingDir>\OPL63.labs\Scheduling\Staff\solution\sched\_staffSolution\Step3.mod



Do not yet copy the synchronization constraint into your **work** copy. First you are going to solve the model and observe the results.

2. Solve the model by right clicking the **Step3** run configuration and selecting **Run this** from the context menu.
3. Look at the **Engine log** and **Statistics** output tabs, and note that the number of variables and constraints treated in the model has increased slightly.
4. The results in the **Solutions** output tab show values for four decision values now, as well as the solution.

### Synchronize simultaneous operations and observe the effects

1. Declare a constraint that synchronizes each task-resource **operation** interval that is present with the associated task's **demand** interval:

```
forall (r in requests, d in demands : d.request == r)
  synchronize(titasks[d], all(o in opers : o.dmd == d)
  tiopers[o]);
```

2. Solve the model by right clicking the **Step3** run configuration and selecting **Run this** from the context menu.
3. Look at the **Engine log** and **Statistics** output tabs. The number of variables and constraints treated in the model has increased significantly, as has the number of fails.
4. Look at the results in the **Solutions** output tab, and note, especially how values for **workers** have changed from the previous solve.
5. Close **Step3.mod**.



## Surrogate constraints

While these constraints completely describe the model, at times it is beneficial to introduce additional constraints that improve the search. Along with treating the resources individually, you can also treat the set of resources of a given type as a **resource pool**. A resource pool can be modeled using a cumulative function expression.

Each resource type has one **cumulFunction** associated with it. Between the start and end of a task, the **cumulFunction** for any required resource is increased by the number of instances of the resource that the task requires, using the **pulse** function.

A constraint that the **cumulFunction** never exceeds the number of resources of the given type is added to the model.



These surrogate constraints on the **cumulFunction** expressions are crucial as they enforce a stronger constraint when the whole set of resources of the tasks is not chosen.

Complete the **Staff Scheduling** workshop by performing the **Add a surrogate constraint to accelerate search** step.

You can perform this lab using the HTML workshop, or by following the instructions in the workbook. The HTML workshop will give you direct access to OPL documentation pages that can help you with the lab.

## Add a surrogate constraint to accelerate search

### Actions

- Declare the cumulative function
- Constrain the cumulative function
- Solve the model and examine the results

### Reference

*cumulFunction*

### Declare the cumulative function

1. Open **Step4.mod** for editing.
2. To model the surrogate constraint on resource usage, a cumulative function expression is created for each resource type. Each **cumulFunction** is modified by a **pulse** function for each demand. The amount of the **pulse** changes the level of the **cumulFunction** by the number of resources of the given type required by the demand:

```
cumulFunction cumuls[r in ResourceTypes] =
    sum (rc in requirements, d in demands : rc.resource == r &&
        d.task.type == rc.task) pulse(titasks[d], rc.quantity);
```

### Constrain the cumulative function

1. You will see that a new intermediate data declaration exists:

```
int levels[rt in ResourceTypes] = sum (r in resources : r.type ==
    rt) 1;
```

This is used to test for the presence of a given resource in a resource type.

2. Write a constraint that requires, when a resource is present in a resource type, that the value of the **cumulFunction** must not exceed the value of **levels**.
3. Compare your results with the contents of  
`<trainingDir>\OPL63.labs\Scheduling\Staff\solution\sched_staffSolution\Step4.mod`

### ***Solve the model and examine the results***

- *Solve the model by right clicking the **Step4** run configuration and selecting **Run this** from the context menu.*
- *If you look at the **Engine log** and **Statistics** output tabs, you will note a dramatic improvement in the number of fails, thanks to the surrogate constraint.*