

SchedulePro Cloud User's Manual

Contents

1	Quick start	3
1.1	Planning intro	3
1.1.1	Landing page	4
1.1.2	Application layout	4
1.1.3	Tables and side panels	6
1.2	Basic scheduling walkthrough.....	8
1.2.1	Creating a schedulable recipe	8
1.2.2	Scheduling the recipe.....	10
1.3	Managing workspaces and scheduling boards	11
1.3.1	When to use single vs. multiple workspaces	11
1.4	Production tracking.....	12
2	Scheduling concepts	16
2.1	Recipes	16
2.2	Resources	17
2.3	Scheduling.....	18
2.4	Task precedence.....	19
2.5	Conflict Resolution	21
2.6	Recipe attributes / changeover matrices	22

1 Quick start

SchedulePro Cloud is a web-based, cloud-native, lightweight APS/MES software system, offering advanced planning and scheduling capabilities as-a-service or via dedicated hosting. The system accounts for task precedence, resource constraints (e.g. equipment, staff, labor), material availability, and job priorities, and can be used for production planning and scheduling, capacity analysis, facility design and debottlenecking. When used for day-to-day production scheduling, schedules can be optionally shared with the plant floor, allowing operators to view scheduled tasks information and provide feedback on actual task execution.

SchedulePro Cloud is architected as a multi-service system, consisting of a number of web front-end applications and back-end services. APS functionality is provided through the Planning application while MES functionality is provided through the Production application. Planning employs a comprehensive, priority-based scheduling algorithm that satisfies all required constraints (when possible) and reports on conflicts when constraints are not met. Schedules can be visualized and customized through various interactive user interfaces, and, optionally, published to Production for execution. Feedback received from Production becomes readily available in Planning and can be reviewed, before it is incorporated into the planned schedule and republished. Production planners develop (and optionally share) their schedules in Planning, while operators log onto Production to view schedule tasks and provide feedback. Powerful equipment occupancy charts allow for visually interacting with the schedule and help to quickly narrow in on major schedule deviations by providing simultaneous display of original, planning, and current schedule versions.

The Planning application can be accessed at <https://planning.scpcloud.intelligen.com> and the Production application at <http://production.scpcloud.intelligen.com>. Before using the system, you need to register for a demo account by clicking the relevant link on the application logon page.

1.1 Planning intro


Planning data are organized in workspaces and scheduling boards. A workspace is the top-most organizational entity; it includes the workspace assets (such as options, materials, resources and recipes) and one or more scheduling boards. Workspaces are isolated from one-another, such that, entities of one workspace cannot be referencing entities of another workspace. Scheduling boards contain schedules created according to the workspace recipes that run on the workspace resources. A scheduling board contains a number of campaigns, each, in turn, consisting of one or more batches, with each batch being an instance (or materialization) of a recipe. Having multiple scheduling boards in the same workspace allows for experimenting with various scheduling scenarios, while maintaining a common set of recipes and resources. As of the writing

of this text, scheduling boards are isolated from one-another, in the sense that resource allocations on one board do not appear on other boards.

1.1.1 Landing page

The landing page of the Planning application offers the following main options:


- Create sample workspace / scheduling board (left-hand side): Allows for creating a pre-populated workspace and scheduling board with data pertaining to the application tutorials or industry-specific examples. The sample workspace includes the required assets (resources and recipes), while the scheduling board includes a list of already scheduled campaigns (which reference the workspace assets).
- Create empty workspace / scheduling board (right-hand side): Create an empty pair of workspace and scheduling board that need to be populated from scratch.



Selecting the sample workspace / scheduling board option displays a list of the available examples. After selecting the desirable example, the system creates the relevant workspace and scheduling board, and schedules one or more scheduling board campaigns, depending on the example contents. It then redirects to the EOC (equipment occupancy) page of the scheduling board, which allows for visualizing and modifying scheduled tasks.

1.1.2 Application layout

The figure below displays the EOC page for Tutorial 1b.



The application layout consists of the top part of the page (areas 1-4), shared by all application pages, and the bottom part (area 5), which is page-specific. The bottom part of the page in the above screenshot corresponds to the equipment occupancy chart including the various related display options and actions.

The 5 areas of the application UI as marked in the previous screenshot are briefly described below:

1. “SCP Cloud Planning” is the application name and remains the same on all pages; clicking the text switches to the landing page. Below the application name is the workspace name (in this example “Tutorial 1b”), and, further down, the application breadcrumbs.
2. “Scheduling-board-1” is the name of the currently selected scheduling board. All 5 navigation items inside area 2 are related to the selected scheduling board and the contents of their corresponding pages will generally differ when a different scheduling board is selected:
 - Home (scheduling board home page): Allows for displaying and editing board properties and campaigns. Also, allows for performing scheduling actions (schedule, reschedule, unschedule, conflict resolution) and modifying scheduled tasks in campaign contents.
 - EOC (equipment occupancy chart page): Displays the scheduling board resource occupancy and allows for modifying the schedule. Bars or groups of bars (depending on the selected “Drag” mode) can be dragged horizontally for time shifting. Individual bars for procedures and operations can be also dragged vertically for resource reallocation. Right-clicking a bar gives access to context-

specific actions, including displaying the side panels of the various entities related to that bar. The buttons on the top-right corner provide access to the scheduling-related actions for the current scheduling board.


- Gantt (operations Gantt page): Displays the scheduled tasks from a process perspective. Right-clicking a bar gives access to the side panel of the bar entity.
 - Line Charts (materials, resources, inventory): Displays consumption/production and inventory charts for the scheduling board.
 - Configuration (include / order resources, scheduling): Allows for editing various scheduling board-related options. “Include / Order resources” includes EOC visibility and ordering options, while “Scheduling” includes options relating to the scheduling algorithm.
3. These 4 navigation items pertain to data and actions of the workspace. Their contents are the same for all scheduling boards within the same workspace:
- Recipes (recipes page): Allows for listing, creating, deleting, and editing recipes.
 - Facilities (facilities page): Allows for listing, creating, deleting, and editing facilities and their resources (such as equipment, staff, labor, and storage units).
 - Materials (materials page): Allows for listing, creating, deleting, and editing materials.
 - Others (options pages, exporting/importing functionality): Allows for managing various options and, also, invoking the external file exporting/importing functionality.
4. Logout option.
5. Page-specific part: In this particular example, the page-specific part corresponds to the “Tutorial 1b” EOC. The interactive EOC allows for dragging tasks horizontally (time shift) and/or vertically (resource reallocation).

1.1.3 Tables and side panels


Tabular pages allow for displaying, adding, deleting, and editing the relevant items as described below:

- New items can be added by entering the item’s name and hitting enter or the “+” icon on the far-right. Newly added items appear at the top of the list until the next list refresh, when they are sorted according to the current sorting options. For entities whose order is important, new items appear at the bottom of the list.
- Items can be deleted by selecting one or more of them and clicking the “Delete” button. Alternatively, the solid down-pointing triangle on the far left of each row displays a context menu, which usually offers the delete option for the particular row.

- Some cell values can be edited in-place by clicking the pencil icon (such as the name of the entity that usually appears in the first column).
- Filtering functionality, available from the “Filter” button, allows for creating rich filtering criteria by selecting the top-level relationship (“and” or “or”) and specifying a number of conditions for individual columns. An extensive array of operators, specific to each column data type, gives great flexibility with achieving even the most demanding filtering criteria.
- Sorting functionality, available from the “Sort” button, allows for hierarchically sorting rows on multiple columns (rows with the same value are subsequently ordered by the next column and so on). Sorting functionality is not present for entities whose order plays an important role in the application (e.g. campaigns or procedures/operations in a recipe).
- Clicking the item name (e.g. “Recipe 1”), usually found in the first column of the table, displays the item’s “side panel”, which includes detailed information on the item properties. The links on the far right of the table (“actions”) switch to another page related to the row item. The special link “Contents” switches to the page that displays the contents of the relevant item.



Side panels consist of one or more tabs (or “pills”) at the top of the panel and one or more groups on each tab. Usually, groups are initially in read-only format, in which a synopsis summarizes the selected values for the group and an “Edit” link, expands the group for editing. When in edit mode, changes can be saved or cancelled by clicking the appropriate button. In case of errors, one or more validations message will appear when trying to save. Some validation rules are hard: the rule applies even when the property is disabled (e.g. a property value must always be numeric even when it is disabled), and some validation rules are soft: the rule does not apply when the property is disabled (e.g. a property value must be greater than zero when enabled, but it can be equal to zero when disabled). Some groups are always expanded, in which case, changes are directly saved to the application database.



On the EOC and Gantt charts, side panels for displaying and editing task information can be brought up by right-clicking a bar and selecting “Properties” for the desired entity in the bar context menu.

1.2 Basic scheduling walkthrough

In order to schedule campaigns on a scheduling board at least one schedulable recipe must be present. At a bare minimum, a schedulable recipe must include one procedure, with at least one piece of equipment in its main equipment pool, and one operation within that procedure. Generally speaking, recipes reference resources included in facilities, so, before making a recipe schedulable, we need to have at least one facility with at least one piece of equipment to be used as main for the sole procedure of the recipe.

1.2.1 Creating a schedulable recipe

If a facility is not already there, a new facility can be created by typing the facility name and hitting “Enter” or the “+” button on the far-right of the table row (as described in the previous section). Clicking the “Equipment” link in a row, switches to the list of equipment for the facility of that row. Here, a new piece of equipment can be added in the usual way. For using the equipment in the main equipment pool of this example, the default options suffice. In any event, equipment properties can be viewed or changed by clicking the equipment name and accessing the relevant side panel.

The next step involves creating and configuring a recipe so that it can be scheduled. A new recipe can be created on the “Recipes” page, by entering the desired name in the recipes table. Then, clicking “Contents” switches to the contents of the new recipe. The “Status” reported in the info box at the top of the page shows that the recipe is invalid (i.e. it cannot be scheduled) because it

contains 1 error. Clicking the error gives more information, namely, the recipe contains no procedures, thus it cannot be scheduled.

A new procedure can be created by entering the procedure name in the table. After creating the procedure, the previous recipe error has been replaced by two new errors: 1) The procedure contains no main equipment, 2) The procedure contains no operations. Clicking the pencil icon in the “Main Equipment Pool” column of the procedures table allows for specifying one or more equipment that can be used for that procedure. Once we select the equipment created previously and hit the “Enter” key or click the green check button to confirm the selection, the errors are back down to 1 (the error about the procedure containing no operations).

In order to create an operation, the procedure contents need to be expanded by clicking the V-shaped icon in the beginning of the row. Entering the operation name adds a new operation with the default duration of 1 h and causes the recipe status to become valid, amounting to a schedulable recipe. The recipe should now look similar to the screenshot below.

SCP Cloud Planning
Tutorial 1b

Scheduling-board-1

Home Eoc Gantt Line Charts Configuration Recipes Facilities Materials Others Logout

Recipes > Recipe

Name Recipe Batch time 1 h Description
Size 1 kg Minimum cycle time 1 h Status Valid Scale Gantt Edit

Path Level ☐ Show Branches ☐ Show Sections ☐ Hide Link Arrows

Procedure Name	Procedure Description	Main Equipment Pool
☐ Procedure		Main Equipment (Facility)


Operation Name	Resources	Duration	Streams	Links
☐ Operation		1 h		S F

Enter operation name in procedure Procedure...

Enter procedure name in recipe Recipe...

The vertical line segment in the “Links” column of the operations table represents the scheduling link timing mode for the operation. When there is no arrow pointing at the start, the operation starts with the start of the batch. If a new operation were to be added to the procedure, its timing link would be automatically set to Finish-to-Start with respect to the existing operation (you may experiment with this at your discretion).


Additional procedure and operation configuration can be done via the corresponding side panels, by clicking the procedure or operation name respectively. Additional scheduling- and resource-related options can be set at the operation level as shown in the screenshot below.



1.2.2 Scheduling the recipe

In order to schedule the recipe, a new campaign needs to be created on the scheduling board “Home” page. Once a campaign is created, we need to associate it with the recipe created previously (it can be done in-place by clicking the pencil icon in the “Recipe” column, or via the “Identification” group in the “General” tab of the campaign side panel). For this example, the number of batches is also edited to 4 from the default value of 1.


Clicking the “Schedule” button inside the info box at the top of the “Home” page, schedules all unscheduled campaigns of the scheduling board (buttons inside the box refer to the entire board). We could achieve the same result by clicking the “Schedule” button above the table after the desired campaigns have been selected from the corresponding check box. Once the campaign is scheduled, we can view the scheduled tasks in tabular form by clicking the “Contents” link for that campaign. The batches can be expanded all the way to their operations. Clicking the name of an entity displays the entity side panel, allowing to view or modify the relevant information.



The EOC page visualizes the scheduling board contents by means of a scheduler Gantt chart. It also allows for interacting with the schedule as briefly described in 1.1.2.

1.3 Managing workspaces and scheduling boards

In addition to creating sample/empty workspaces and scheduling boards as described in section 1.1.1, the landing page allows for adding new scheduling boards in the selected workspace (if a workspace is currently selected) and switching to another workspace/scheduling board (if one is available). The relevant link at the bottom right of the page allows for getting to the “Manage workspaces & scheduling boards” page displayed in the screenshot below.



A number of options are available on that page:

- A new workspace can be created; initially the workspace has no scheduling boards and, thus, cannot be selected.
- A new scheduling board can be added to an existing workspace.
- An existing workspace/scheduling board can be entered (from the “Go to” link).
- Production mode can be initiated/terminated for a scheduling board (more on that in the next section).
- Workspaces and scheduling boards can be deleted, and their names and descriptions can be in-place edited.

1.3.1 When to use single vs. multiple workspaces

An organization may have to decide if its production modelling/scheduling needs can be satisfied by using a single workspace with one or more scheduling boards or if multiple workspaces are required. Different organizations may have different views as to how organization data should be structured, so there is no definitive answer to that question. However, some general guidelines can be considered.

For small organizations, with a small number of production lines that share most of their resources, a single workspace would be the most practical solution. A main scheduling board can be used for the core planning and/or production tracking activities, while additional scheduling boards can be used to explore alternative scenarios in planning mode (e.g. capacity analysis, debottlenecking etc.), so as not to interfere with the main schedule.


For medium-sized organizations with several production lines it would probably still be more practical to keep everything in the same workspace, as most assets (options, materials, resources, recipe) will be mostly likely shared across the organization. Groups of closely-related production lines (i.e. those sharing resources) would be scheduled on the same board. In order to track production, production features would be enabled for each of these boards. Again, additional scheduling boards could be used to explore alternative scenarios in planning mode.

Finally, larger, possibly multi-site, organizations may opt for employing multiple workspaces to fulfil their planning and scheduling needs. As workspaces are completely isolated, sharing data across workspaces (e.g. options such as outage templates or attention codes) will need to be done manually. The good news is that these options can be exported from one workspace and imported to other workspaces, so maintenance of shared data is largely painless. Each workspace may contain one or more scheduling boards as described in the previous paragraphs.

1.4 Production tracking

Enabling production tracking for a scheduling board can be done by clicking the “Initiate” button for that board on the “Manage workspaces & scheduling boards page”, as shown in the previous section. Production-enabled scheduling boards become available for selection on the landing page of the Production application. At any given moment, Production is associated with a single scheduling board, the contents of which can be displayed and updated (in the context of the Production application, workspaces are not relevant).

In Planning, the EOC page for a production-enabled scheduling board includes a “Planning” vs. “Tracking” toggle at the top-center of the page and a “Sync” button at the far-right of the scheduling-related. The “Sync” link allows for syncing planning data with production. A synced scheduling board defines a “Restricted horizon” during which Planning cannot modify scheduled tasks: these tasks are either in the past or in the immediate future (production may have already started or may be in preparation), so they can only be changed from Production. The restricted horizon appears as a shaded area of the chart and its duration into the future is configurable.



The Production application offers a couple of highly configurable views: the graphical “EOC” (of similar appearance but different functionality as the Planning EOC) and the tabular “Operations”. Unlike Planning EOC, Production EOC does not allow for mass modifications to the schedule, so dragging tasks is disabled. Dragging and right-clicking are generally avoided as the Production application needs to be touch-screen-friendly for use on plant floor portable devices. Clicking a bar in the Production EOC allows operators to update the task with execution information.

Operation "Mix Thoroughly"

Info (Attention: 06/19/2026 09:00 AM)

Campaign	Campaign-1	Main Equipment	MT-1
Batch	Campaign-1-0	Last updated by	-
Procedure	Prepare Solution	Attention codes	-
Comments	-		

Timing

New Status	Completed	Duration	2 h (2 h)
Start	06/19/2026 10:00 AM	End	06/19/2026 12:00 PM
Comment	Attention code Select...		

Auxiliary Equipment


Auxiliary Equipment	-
Planning Auxiliary Equipment	-
Original Auxiliary Equipment	-

Staff


Staff	-
Planning Staff	-
Original Staff	-

Operation timing information can be updated by clicking the “Edit” link in the “Timing” group as shown in the previous screenshot. In Tutorial 1b, delaying the start time by 1 hour and increasing the duration also by 1 hour for the “Mix Thoroughly” operation of the first batch, creates a gap

between that operation and its previous operation, while pushing the rest of the batch forward by 2 hours. Delaying the first batch by two hours, creates conflicts with the second batch on the “CIP Skid” equipment appearing at the top of the chart. As can be seen in the screenshot below, the effects of production updates become immediately visible in the Production application.




At the same time, Planning still displays the “planned” version of the schedule, that is, before the update was entered. However, a prominent link in red color notifies of 1 pending production update. Switching the EOC toggle at the center of the page to “Tracking” and enabling the “Planning baseline” option, displays the planning vs. tracking versions of the schedule, juxtaposing the two versions and allowing for easily spotting schedule deviations.



Clicking the “1 Update” link displays information about the update: the operation it refers to and the new production information for that operation. Text in red highlights the properties whose values have changed compared to the planned schedule. Clicking the “Apply all” button on the “Production Updates” side panel incorporates the updated information into the scheduling board. Doing so, potentially creates scheduling conflicts, the same as those that appeared in Production when the updates were first entered. Switching back to “Planning” mode from the EOC toggle, allows for reviewing and resolving the conflicts as appropriate. Finally, the updated schedule can be published again to Production by clicking the “Sync” button.

Day-to-day scheduling and production tracking involves continuously repeating the updating cycle described above, in intervals that suit the organization modus operandi. Some organizations may opt for updating production at the end of each shift, while others may choose to do so more frequently (or even less frequently, especially when outside regular working hours).

With every cycle, a new schedule version is created, one that incorporates the production updates and the corresponding planning adjustments up to that point. SchedulePro Cloud archives all these schedule versions in a document database and makes them available via a web API, facilitating historical data analysis. Additionally, it provides easy access to information related to the following two older versions of the schedule: original (the state of the batch when it was originally synced with production), planning (the state of the batch at last sync). (The main tracking schedule corresponds to the planning version including updates that haven’t been incorporated yet). The following screenshot displays the production EOC with visible original and planning baselines (the bottom baseline corresponds to original information while the top baseline corresponds to planning information).



2 Scheduling concepts

As briefly described in the “Quick Start” chapter, SchedulePro Cloud uses the concept of a recipe to model a production process in terms of its two main characteristics: timing requirements (durations and precedence constraints) and resource requirements (resource assignments). Scheduling a recipe involves scheduling an associated campaign of one or more batches, where each batch is populated as an instance of the recipe at scheduling time. It is worth noting that recipes do not specify absolute timings, but only timings relative to the process start; it is only at scheduling time that those relative times in the recipe are converted to absolute timings in the batch. In the face of contention for resources, recipes provide flexibility in two main directions:

- Resource assignment: tasks may be associated with pools of alternative resources (as opposed to exact resources), only a subset of which are used in the batch at a given time.
- Timing: tasks may be configured so that they can be delayed or paused until resources become available.

2.1 Recipes

Recipes are organized hierarchically into branches, sections, procedures and operations. Branches and sections are mostly used for backward compatibility and grouping purposes, so they are of no importance to scheduling. By default, they are omitted for new recipes. The core of the process modelling information is held in procedures and operations.

Procedures represent distinct manufacturing steps that utilize one primary piece of equipment for its entire duration (called “main equipment”). For example, a mixing procedure might utilize a blending tank. Procedures consist of operations that describe the distinct sub-steps in a procedure. Main equipment is the only resource associated with procedures. Additional resources, such as equipment (called “auxiliary equipment”), staff, labor, and materials, can be specified at the operation level. SchedulePro Cloud does not distinguish between main and auxiliary equipment, allowing for the same equipment to be used as main or auxiliary according to modelling needs.

Procedures do not hold any timing information; instead, timing information is determined by their operations: the procedure start is the earliest operation start and the procedure end the latest operation end. The main equipment is considered occupied for the entire duration of the procedure even during gaps between operations or gaps (called breaks) within operations.

The relative timing of operations is determined by their duration and their scheduling timing links. Operation durations can be specified as follows:

- Fixed: Set by the user to a constant value.
- Rate-based: Calculated based on a rate, depending on the processed amount.

- Schedule-dependent: Equal to the duration of another operation or series of operations.
- Changeover-dependent: Drawn from a changeover matrix that defines transitions between products.

Operation scheduling timing can be set according to one of the following options:

- Start relative to the start of batch (batch start).
- Start relative to the start of another operation (Start-to-start, SS).
- Start relative to the end of another operation (Finish-to-start, FS).
- End relative to the start of another operation (Start-to-finish, SF).
- End relative to the end of another operation (Finish-to-finish, FF).

The above specifications are used as a starting point, on top of which a fixed or flexible shift time may be applied. A fixed shift is a delay (or advance) that is part of the process and, thus, always applied. A flexible shift is a way to specify that it is acceptable for certain tasks to start later (or earlier) than their prescribed time and is only applied if the operation needs to wait for a resource to become available. An operation can also be declared as 'interruptible' meaning that it can pause and resume operation when the resources it needs become again available.

2.2 Resources

Resources in SchedulePro Cloud represent the physical items that are necessary to schedule a recipe (i.e. create a batch) and belong to the following categories:

- Materials: Specify the various materials referenced in the recipe, without including actual availability information. Availability is determined by the available inventory of a material in special tracking entities called storage units. Materials are tracked by means of input and output streams, which are specified at the operation level. Each stream contains one or more materials, whose amount is specified as a fraction of the overall stream amount in the stream composition. Streams can be drawn from or deposited to a storage unit.
- Unit resources (equipment, staff): Equipment represents the manufacturing equipment needed to carry out a specific process step, for example a tank, a reactor or a filling machine. From a scheduling point of view, an equipment item is a non-consumable reusable resource. Equipment can be assigned to procedures and operations as main and auxiliary equipment respectively. Staff, on the other hand, which represent individuals, can only be assigned to operations. All unit resources are assigned in pools of alternatives. For example, a mixing step that can be carried out in Tank-1 or Tank-2 would include both tanks in the pool. SchedulePro Cloud will pick the first available tank for a given batch.
- Continuous resources (labor): Labor represents the person effort required for certain process steps that is not tied to any specific individual (contrary to staff defined above). It is usually measured in persons.

Scheduled entities (operations, procedures, batches) might need to be delayed until the required resources become available. All resources (except materials) are organized into facilities. Facilities and their resources are equipped with calendars where their availability can be declared (e.g. outages, downtimes etc.). Resources act as constraints that the scheduler tries to satisfy, otherwise conflicts are generated. SchedulePro attempts to schedule in a way that does not exceed the available resources.

2.3 Scheduling

Scheduling starts with creating one or more campaigns on a scheduling board. Each campaign consists of one or more batches, with each batch being an instantiation of the recipe. Additionally, campaigns contain information related to timing/sequencing (start time or due time, independent or related to some other campaign), recipe scaling (scale factor or target batch size), cycling (time interval between the start of two consecutive batches that can be calculated or manually set). Campaigns may also contain pre-production and post-production steps to represent time spent setting up or cleaning out the equipment. These non-production steps are treated as regular operations that may consume materials and resources.

When a campaign is first created, its batches are empty; at scheduling time, batches are populated with procedures and operations created as instances of the procedures and operations of the recipe. In order to distinguish between recipe and batch tasks, we sometimes refer to batch tasks as scheduled procedures and operations, though we usually omit the characterization when the context is clearly related to the schedule rather than the recipe.

At a high level, SchedulePro Cloud schedules a campaign as follows:

- Creates scheduled procedures and operations for every batch according to the recipe. Scheduling timing and duration relationships are created for the scheduled tasks according to the recipe. Batch start times are set according to the campaign cycle time, which is by default automatically calculated, but can also be set manually.
- Assigns default unit resources to newly created procedures and operations. Default resources are assigned per batch in a round-robin fashion from the relevant resource pools. For example, a recipe procedure with a main equipment pool of two items, will be assigned the first item in the first batch, the second item in the second batch, again the first item in the third batch etc.
- Sets the timing and duration information for operations in each batch according to the relationships created in the first step.
- Identifies conflicts within or across batches that may have been created in the previous steps. Conflicts are individual constraint violations related to a particular resource for a specified time interval.

- Resolves individual conflicts. After resolving each conflict, the conflicts are again identified as resolving one conflict may have affected other conflicts in the schedule.

Setting timing and duration information for a batch is performed according to our proprietary dependency algorithm. The algorithm constructs a graph whose nodes correspond to the start and duration of each operation in the batch and whose directed edges correspond to dependencies between operation starts and durations. A topological sort of the graph yields the order in which the start and duration of operations within a batch need to be determined. In case of cycles, the topological sort reports on the first cycle that was encountered (a recipe with cycles is not schedulable).

SchedulePro Cloud supports the following types of constraints/conflicts:

- Main equipment overallocation: The maximum number of allowed parallel uses for some main equipment has been exceeded (usually this value is 1).
- Main equipment outage: Some main equipment is used during an outage.
- Main-to-main equipment incompatibility: The main equipment for a procedure is incompatible with the main equipment of its compatibility master procedure. This constraint is used to model suites of equipment within which only a subset of the equipment in the pool is available.
- Auxiliary equipment overallocation: The maximum number of allowed parallel uses for some auxiliary equipment has been exceeded (usually this value is 1).
- Auxiliary equipment outage: Some auxiliary equipment is used during an outage.
- Main-to-auxiliary equipment incompatibility: The auxiliary equipment for an operation is incompatible with the main equipment of the relevant procedure. Similarly to the main-to-main equipment compatibility, this constraint is used to limit the auxiliary equipment in the pool to only the subset that is compatible with the main equipment selection.
- Staff overallocation: The maximum number of allowed parallel used for some staff has been exceeded (usually this value is 1).
- Staff outage: Some staff is engaged in a task during an outage.
- Labor overallocation: The overall availability for some labor type has been exceeded.
- Campaign timing violation: The timing constraints of a campaign with respect to its master campaign has been violated.
- Storage unit maximum rate violation: The maximum incoming or outgoing rate for a storage unit has been violated.
- Storage unit inventory: The inventory limits for a storage units (low or high) have been violated.

2.4 Task precedence

Resource allocation in SchedulePro Cloud is based on task precedence: tasks of higher precedence have a higher priority when claiming resources than tasks of lower precedence. By that token, when multiple tasks are simultaneously using a resource such that the resource capacity is exceeded (overallocation conflict), it is the lower-precedence task that is considered responsible and, thus, it is the one that needs to be moved for the conflict to be resolved. Other types of conflicts, such as those related to outages, compatibility or inventory, only involve a single task, so task precedence is not relevant. Task precedence is assigned hierarchically, starting at the campaign level, to the batch level, to the procedures and operations within a batch as described in the following paragraphs.

Campaign precedence is determined by the order in which campaigns appear on the scheduling board home page (top campaigns have higher precedence than bottom campaigns). Similarly, the order in which batches appear within each campaign, determines the batch precedence (top batches have higher precedence than bottom batches). Exceptionally, batch precedence for backward scheduled campaigns (i.e. campaigns with a due time) is assigned in reverse order (bottom batches have higher precedence than top batches).

Precedence assignment to individual procedures and operations within a batch is performed according to our proprietary dependency algorithm described above. The sorted node list, produced by the dependency algorithm, is used to create a list of batch operations in dependency order. Assigning precedence to procedures and operations involves iterating over the ordered operations list twice: at the first pass, procedure precedence is assigned in the order in which procedures are encountered; at the second pass operation precedence is assigned. The precedence of procedures or operations that are not on the same dependency path is determined by the order in which they appear on the recipe Contents page.

Precedence value is assigned sequentially for procedures and operations, such that all procedures within a batch end up having higher precedence than the batch operations. Since the only resource allocated directly to a procedure is its main equipment, the above results in procedures having a higher priority in claiming main equipment than operations do auxiliary equipment. Priority among procedures or operations is determined by their dependency order.

In short, precedence assignment follows the rules below:

- Campaign precedence follows the order in which campaigns are listed on the scheduling board Home page.
- Batch precedence within a campaign follows the order in which batches are listed on the campaign Contents page. For a backward-schedule campaign, the reverse order is followed.

- All main equipment assignments within a batch have higher precedence than auxiliary equipment assignments within the same batch.
- Precedence of main equipment assignments to procedures within a batch follows the scheduling link dependency order.
- Precedence of other resources to operations within a batch follows the scheduling link dependency order.
- The precedence of procedures or operations not on the same dependency path is determined by their order on the recipe Contents page.

TODO: Resource conflict precedence

2.5 Conflict Resolution

Conflicts are identified by building the utilization profile for the various resources involved in the schedule. When all conflicts need to be identified (e.g. for display purposes), the utilization of resources is calculated for the entire schedule, until all tasks have been accounted for. Conflicts are identified and listed in precedence order of the conflicting task. That is, conflicts appearing higher on the list are related to tasks of higher precedence as compared to conflicts in lower positions.

When conflicts need to be resolved (either during the initial scheduling phase or at user's request), the algorithm follows the steps below:

1. Calculates the utilization of resources in task precedence order, until the first conflict is identified.
2. Tries to resolve the conflict. If the conflict cannot be resolved, it is marked as an unresolvable conflict.
3. Recalculates the utilization of resources in task precedence order, until the first non-unresolvable conflict is identified.
4. Goes back to step 2.

At a high level, resolving a single conflict involves trying the following options in that order:

1. Search for an alternative resource to move the task to. This is only possible for unit resource conflicts (e.g. equipment and staff). If an alternative resource is found, the task is moved to that resource and conflict resolution exits. Otherwise, conflict resolution proceeds to the next option.
2. If the operation is interruptible, try adding a break for the conflict duration. Since main equipment is considered occupied even during operation breaks, adding a break is not a viable resolution strategy for main equipment conflicts. If a break can resolve the conflict,

the break is added to the operation and conflict resolution exits. Otherwise, conflict resolution proceeds to the next option.

3. If the task is configured for flexible shifting with respect to its scheduling link, try adding a flexible shift for resolving the conflict. In the case of a main equipment conflict, all operations within the conflicting procedure will need to be shifted appropriately. If a flexible shift can resolve the conflict, the flexible shift is added and conflict resolution exits. Otherwise, conflict resolution proceeds to the next option.
4. Try shifting the entire batch in order to resolve the conflict. Except for some special cases, batch shifting is only good for inter-batch conflicts (conflicts that involve different batches). If a batch shift can resolve the conflict, the batch is shifted and conflict resolution exits. Otherwise, the conflict is deemed unresolvable.

Notice how the conflict resolution options go, generally speaking, from least to most disruptive. The algorithm is trying to resolve conflicts starting with the options that would have the smallest overall impact on the schedule and gradually moves to the ones whose impact would be more extensive.

2.6 Recipe attributes / changeover matrices

It is often the case that different recipes share scheduling options such as equipment processing rates or setup/changeover times. Recipe attributes enable grouping recipes so that common options can be configured for entire recipe groups, rather than for individual recipes. Each recipe attribute defines a list of possible values and each recipe can be assigned at most one value per available attribute. Assigning values from multiple attributes, allows for creating recipe groups based on different characteristics, which are orthogonal with one another.

A typical example is the packing size for some final product. Often, the packing line processing rate only depends on the packing size, rather than the actual material that goes into the package. For example, recipes whose final products go into a predetermined set of packages for 200g, 500g, and 1kg can be grouped using a “Packing size” attribute that includes 3 values, one for each of the aforementioned packing sizes. The packing line processing rate can then be specified in terms of the value of the packing size attribute, so that, packaging for all recipes with the same value is done at the same rate.

Switching from one packing size to another may incur some setup time on the packing line. Also, the setup time may differ depending on the previous vs the next packing size that is processed. This situation can be modelled with the help of changeover matrices. Changeover matrices are associated with recipe attributes and determine the setup/changeover times for transitions between the values of the attribute they are associated to. Depending on the modelling requirements, there may be multiple changeover matrices per recipe attribute.