



Where are you going to store it?

Take a perfectly spherical pink elephant whose mass is negligible ...

Management of intermediate storage is not a popular topic of discussion in sales literature for process scheduling systems. Brochures and web sites tout “state of the art algorithms”, “artificial intelligence”, “available to promise”, “just in time scheduling”, and a dozen other popular buzzwords creating an aura of highly sophisticated scheduling technology. Invariably, these same sources never say a word about how they deal with intermediate storage. Why the deafening silence on this subject? Is the storage of process intermediates an unimportant detail in process scheduling, or are other vendors just ignoring the problem? In the following section we explain why generating schedules in which the storage pattern of intermediate materials is both fully specified and feasible is the holy grail of process scheduling. So much so, that a brochure touting “state of the art optimization” in a system which cannot handle intermediate storage would perhaps better be titled, “Take a Perfectly Spherical Pink Elephant, Whose Mass Is Negligible...”

Consider a simple two stage process, with materials A and B produced in large batches then packed on packing lines. Certainly the packing lines cannot consume A or B before the material is created. Many scheduling systems on the market can help the user to create schedules which are feasible in this respect. That is, before a packing task which consumes 1,000 lbs of A is scheduled, there will be sufficient making tasks for A in the past to provide the material. This is usually evident by the lack any negative numbers on the inventory plot for material A, though many systems require the user to perform manual manipulations to remove negative inventories. But real life scheduling is a little more complicated than that. Not only must the inventory plots for A and B be always non-negative, but for any time in which there is a positive inventory, there must be a vessel available to hold that material. Scheduling without handling this constraint is like designing an airplane without considering gravity. It is really easy, and it won't fly.

We have yet to encounter a process with separate infinitely large silos or tanks for each and every intermediate product. How realistic is it then to use a scheduling system which assumes that every intermediate can be stored in whatever quantity is desired? Many scheduling systems detect inventory levels above the allowed storage quantity and expect the user to manipulate the schedule to remove these violations. How much “optimization” is a system doing if the user must fiddle with the solution just to achieve feasibility? Never mind all of the hype about “available to promise”, the educated customer will ask “can the system properly handle simple material balance constraints that every sophomore chemical engineering student knows are the foundation of process operation?”



Even after the user has moved tasks around to get the maximum inventories to the “right” levels, just what is the right level? Many processes have dozens of intermediate materials which are storable in two or three tanks. Just what is the maximum inventory level for an intermediate in such a process? Is it the capacity of one tank, or all three? Certainly one could conceivably have two of the three tanks used to store a single material. A bigger question is, “Even if all the inventories are below the maximum level, suppose I have positive inventories of six intermediates and only three tanks?” This is a very common situation; in fact it is the rule rather than the exception. A system which can only “help” the user to keep inventories of the individual intermediates below the capacity of one of the tanks, does nothing to enforce the real world constraint that the number of intermediates with positive inventory cannot exceed the number of tanks. Sure you can make it, but where are you going to put it?

A fundamental law of scheduling is that material that is in inventory must actually be somewhere. Many scheduling systems were developed for use in manufacture of items like shoes, or engine blocks. Because such materials can be stacked side by side in a warehouse, it is easier to ignore storage considerations in such cases. But, if you are dealing with liquids, or unpackaged solids, e.g. cereal or crystals, you don’t have a feasible schedule unless you know where every pound of material is stored at every time. VirtECS™ provides this knowledge.

VirtECS™ generates only feasible schedules and it knows (and can tell you) exactly how much of every material you have at any time point on the schedule, and in which vessel that material resides. If you have fourteen intermediates with only three silos, our schedule will have no more than three intermediates present at any one time and we will know which silos are used for which intermediates. When one silo is emptied, then and only then will VirtECS use it to store another material. VirtECS also handles “process vessel only” storage, where materials are storable only in the vessels which made them. In such a case, VirtECS will not start another task on that vessel until all of the material from the previous batch has been used in downstream tasks.

🌀 Scheduling Pitfall: If you don’t have infinite dedicated storage of materials, if you have process vessel only storage, or shared storage vessels, be wary when evaluating a scheduling package to confirm it can automatically handle your storage constraints. 🌀

