Optimising the wine supply chain

M. Michalewicz[1], Z. Michalewicz[1,2], R. Spitty[3]

[1]SolveIT Software, Pty Ltd, Level 1, 99 Frome Street, Adelaide, SA 5000, Australia. [2]School of Computer Science, The University of Adelaide, Adelaide, SA 5005, Australia. [3]SolveIT Software, Pty Ltd, Level 2, 198 Harbour Esplanade, Docklands, Vic 3008, Australia
Corresponding author's email: mm@solveitsoftware.com

Abstract

Most production-based businesses in general, and wineries in particular, have been under enormous pressure to improve their top-line growth and bottom-line savings during the last few years. As a result, many companies are turning to systems and technologies that can help optimise supply chain activities and improving short – and long-term demand forecasting. Given the inherent complexities of planning and scheduling an agricultural supply chain, many new solving methods (e.g. ant systems, evolutionary algorithms, fuzzy systems, genetic algorithms, neural networks, rough sets, swarm intelligence, simulated annealing, and tabu search – collectively known as 'Computational Intelligence' methods – have been introduced into software applications to help manage and optimise this complexity.

Introduction

This decade has witnessed the emergence of systems and technologies that can help wineries to optimise various business processes, predict and deal with unexpected events, and address key operational issues, such as:

- creating optimal harvest schedules that can accommodate last-minute changes;
- maximising the utilisation of crushers, pressers, and fermentors;
- optimising tank farm transfers and activities;
- dealing with sudden changes (e.g. delayed transport, demand spikes, equipment failure, extreme weather conditions);
- improving resource allocation under risk and uncertainty;
- minimising transportation and other logistics costs; and
- improving the accuracy of demand forecasts.

Computational Intelligence methods are very well suited for powering software applications for addressing these operational issues; we discuss them briefly in the following section.

Computational intelligence

Computational Intelligence is considered an alternative to classical artificial intelligence and it relies on heuristic algorithms (such as in fuzzy systems, neural networks and evolutionary computation). In addition, Computational Intelligence also embraces techniques such as swarm intelligence, fractals and chaos theory, artificial immune systems, and others. Computational Intelligence techniques often combine elements of learning, adaptation, evolution and fuzzy logic to create programs that are, in some sense, intelligent.

An interesting question, which is being raised from time to time, asks for guidance on the types of problems for which Computational Intelligence methods are more appropriate than, say, standard Operation Research methods. From our perspective, the best answer to this question is given in a single word: *complexity*. Let us explain. Real-world problems are usually difficult to solve for several reasons. These reasons include the following:

- the number of possible solutions is so vast that they constrain an extensive search for the best answer;
- the evaluation function that describes the quality of any proposed solution is noisy or varies with time, thereby requiring not just a single solution but an entire series of solutions; and
- possible solutions are so heavily constrained that constructing even one feasible answer is difficult, let alone searching for an optimum solution.

Naturally, this list could be extended to include many other possible obstacles. For example, we could include noise associated with our observations and measurements, uncertainly about given information, and the difficulties posed by problems that have multiple and possibly conflicting objectives (which may require a set of solutions rather than a single solution). All these reasons are just various aspects of the complexity of the problem.

Note that every time we solve a problem, we must realise that we are in reality only finding the solution to a *model* of the problem. All models are a simplification of the real world – otherwise they would be as complex and unwieldy as the natural setting itself. Thus, the process of problem solving consists of two separate general steps: (i) creating a model of the problem; and (ii) using that model to generate a solution:

Problem → Model → Solution

Note that the 'solution' is only a solution in terms of the model. If our model has a high degree of fidelity, we can have more confidence that our solution will be meaningful. In contrast, if the model has too many unfulfilled assumptions and rough approximations, the solution might be meaningless, or worse. So, in solving real-world problems there are at least two ways to proceed:

- we can try to simplify the model so that traditional methods might return better answers; or
- we can keep the model with all its complexities, and use non-traditional approaches, to find a near-optimum solution.

In other words, the more complex the problem (e.g. size of the search space, evaluation function, noise, constraints), the more appropriate it is to use a non-traditional method, for example, the

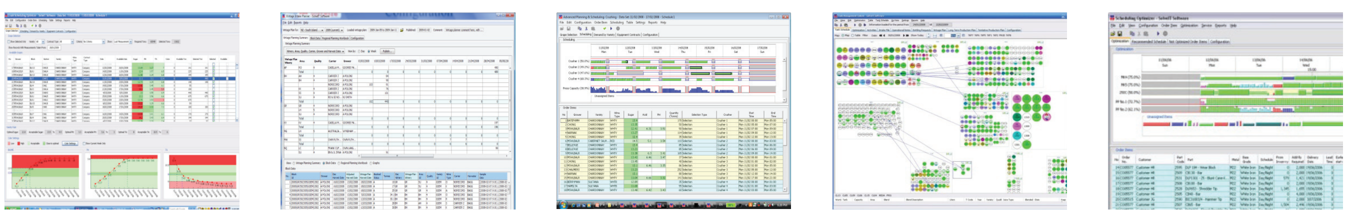

Maturity Models — Vintage Planning — Crushing — Tank Farm — Bottling

Figure 1. SolveIT Software's software applications for optimising the wine supply chain

Computational Intelligence method. Anyway, it is difficult to obtain a precise solution to a problem because we either have to approximate a model or approximate the solution. And a large volume of experimental evidence shows that this latter approach can often be used to practical advantage (see, for example, www.SolveITSoftware.com).

Many Computational Intelligence methods (e.g. ant systems, evolutionary algorithms, fuzzy systems, genetic algorithms, neural networks, rough sets, swarm intelligence, simulated annealing, and tabu search) have already been incorporated into software applications that handle levels of supply chain complexity that is unapproachable by traditional methods.

Wine supply chain

To address many wine production challenges present in different parts of the wine supply chain, SolveIT Software has developed a suite of software applications that can optimise the end-to-end wine supply chain. These software applications are based on the paradigms of Computational Intelligence. As seen in Figure 1, these software applications include predictive modelling for grape maturity (using weather forecasts and readings on Baumé, pH, and TA), vintage planning, crush scheduling, tank farm optimisation, and bottling-line sequencing. When deployed together, these applications can optimise all planning and scheduling activities across a winery's entire supply chain.

The next section discusses these modules briefly.

Maturity models

Vintage intake plans are heavily dependent on the prediction of expected grape maturity dates. It is possible to export the prediction dates from some external system that functions as a black box providing only one date when it believes the harvesting should occur. However, limited visibility into the prediction process often prompts requests to revisit the prediction functionality of this process. So the maturity models deploy a new prediction module that provides improved prediction dates and visibility of the prediction-calculation provided.

Grape maturity can be defined as the physiological age of the berry on the vine. It is important to define the optimal grape maturity for wine production and to develop clear chemical or biochemical traits that can be used to define the peak of ripeness. The definition of optimal maturity will vary depending upon the style of wine being made; the working definition of quality; varietal; rootstock; site; interaction of varietal, rootstock and site; seasonal specific factors; viticultural practices; and downstream processing events and goals. If a clear descriptive analysis of the quality target exists, then the time of harvest can be optimised to meet those goals. Several grape and cluster characteristics have been used to assess ripeness (e.g. sugar, pH, acidity, berry metabolites, berry proteins, taste). There are, of course, other non-compositional factors that influence the decision to harvest, including labour availability; seasonal changes such as rainfall; heat waves; tank space limitations; and other factors beyond the winemaker's control.

A 'black box' prediction approach provides no audit capability for the user making it difficult to detect and promptly address issues related to accuracy of prediction. These factors can easily cause errors in forecasting maturity dates to go unnoticed and unrectified for prolonged periods of time. Decisions on when to book certain grapes for harvesting and crushing are relying heavily on the experience of the personnel involved in the process, and might result in a non-optimal allocation of harvesting and crushing resources. Each of these situations could result in higher costs for harvesting, transportation, and crushing, and reduction in grape quality.

SolveIT Software's software applications for optimising the wine supply chain is depicted in Figure 2.

Vintage intake planning

Vintage intake planning manages the grape supply intake from the 'vineyard to the weighbridge'. The functionality of this module supports the creation and maintenance of vintage intake plans that satisfy capacity constraints and facilitate the harvesting of grapes during periods of time when the quality is the highest. SolveIT Software's Vintage Intake Planner has been designed around the following key components in the efficient planning of grape supply intake:

- Will the block/sub-block be available for harvesting during the optimum date range period, taking the spray diary into account? Instead of assuming that each block is available for harvesting when required, the Vintage Intake Planner receives information that determines the availability to harvest. Example of the information required is the spray diary for each block. The grower currently supplies information on spray at a block/sub-block level to satisfy the compliance requirement, but also to report on the periods when the block/sub-block is under spray and not available for harvesting.
- What is the best optimal/balanced grape supply intake plan taking into account the predicted optimum quality data range, block/sub-block availability for picking/harvesting constraints, and winery constraints? If a Vintage Intake Planner provides a plan based just on the maturity dates provided, such plan would be unbalanced, as it does not take into account other constraints such as crushing/pressing capacity, logistics, etc. Thus it is important to include all the known constraints and information available at the time and provide a recommended optimal and balanced vintage intake plan. The recommended plan is available for the business users to make manual adjustments if required, save and submit the plan.

Crushing

Crushers are used to process wine grapes and are often connected to different types of pressing machines. The optimal processing capacity of the crushing machines is about 40–45 tonnes per hour. However, but if necessary, it may be increased to 60–80 tonnes per hour. The most important limiting factor is the capacity of the pressing machines and fermentation containers. The processing capacity for the pressing machines ranges from 4 to 12 tonnes per hour depending on the type of grapes. It is important to generate optimal schedules for all crushers over some time horizon. However, the generated weekly schedule may incur frequent changes due to contractual influences, weather conditions, seasonal influences, and daily production variances. When the changes occur, SolveIT Software's Crush

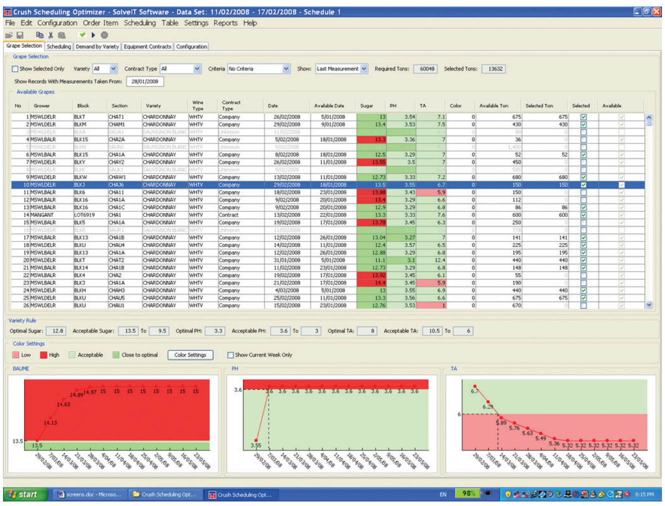

Figure 2. SolveIT Software's predictive model for grape maturity

Scheduler (Figure 3) re-optimises and generates alternative schedules to fill available capacity.

Also, a variety of constraints are present in this part of the wine supply chain, including:

- constraints in time (e.g. not processing fruit during the hottest part of the day);
- constraints in the throughput of presses;
- constraints in the throughput of fermentation vessels;
- the throughput of trucks via the crusher to be a continuous flow;
- scheduled repairs and maintenance of equipment;
- scheduled changeover and clean up (white to red, or lower grade grape-to-higher grade grape); and
- special demand-fulfilling variety shortages to address meeting capacity needs.

Tank farm

Wineries make daily decisions on the processing and storage of wines and juices in their tank farms, with major wineries having several hundred tanks that have differing capacities and attributes. These tanks may be insulated or refrigerated, for instance, and could have an agitator. Some of these tanks might be used to store juice after the grapes are crushed, fermented, and pressed, while others might be primarily designated for blending. Different types of juices may also require specific tank attributes for processing, such as refrigeration jackets or agitators.

The process of planning and scheduling wine and juice transfers on the tank farm is far from trivial, as wine juices that need to be blended later should be kept together in nearby tanks. This ensures lower operational expenses and smoother process flow for the blending process. Second, the tanks should be close to their optimal filled capacity. This would eliminate 'half-empty' tanks and maximise the number of empty tanks available for incoming juice. Third, the allocation of tanks should be flexible enough to accommodate last minute changes due to various unexpected events (e.g. changes in volume, demand, quality).

Until recently, there has been a lack of software tools available that can manage and optimise tank farm activities, and as a result, most wineries have used a whiteboard approach which provides very little forward-looking visibility and does not allow for cost – or quality-optimised decision making. Substantial savings in spillage, labour and electricity use, throughput, and tank utilisation can be achieved through planning and scheduling optimisation.

SolveIT Software's Tank Farm Optimiser (Figure 4) has been developed to replace the whiteboard approach for planning and scheduling tank farm activities. The system generates optimised tank farm plans for the entire vintage, given the physical constraints of the tank farm (e.g. tank attributes, transfer paths), incoming juice, bottling plan, and any user defined business rules. The planning horizon is in the order of 12 months, with daily or weekly 're-optimisation' as changes occur or new information becomes available. Figure 4 shows a screen capture of the 'Map View' panel of the Tank Farm Optimiser, which provides better visibility into the tank farm, warnings if there are conflicting operations scheduled on the same tank on different days, and eliminates the need for using whiteboards to manage the tank farm.

To decrease the overall cost of wine production on the tank farm, the Tank Farm Optimiser uses projected grape intake and bottling requirements to create a production schedule that minimises the number and distance of transfers between tanks, as well as the number of litres being transferred. The Tank Farm Optimiser used advanced non-linear optimisation algorithms (such as genetic algorithms) that consider the following key objectives when creating a production schedule:

- minimisation of pack-up and pack-down transfers;
- minimisation of changeovers from red to white in transfer paths and tanks;
- minimisation of transfer distance;
- minimisation of the number of litres transferred; and
- minimisation of ullage.

Additional objectives, such as environmental factors and quality issues, can also be added to the system to allow for trade-off analysis – these are discussed later in the paper.

The user has full control over the optimisation process by being able to lock in manual decisions, set the business rules and constraints, and re-optimise after making changes. The Tank Farm Optimiser also provides elaborate reporting on a number of levels to suit different stakeholders; for example, the daily bottling programs for execution, a report on particular wine blends, or a report on expected production efficiency.

The Tank Farm Optimiser takes into account the constraints and complexities of wine production to provide users with optimal scheduling of tank operations, environmental impact analysis, and reporting. Through the optimisation, advance planning, and visibility functionalities provided by SolveIT Software's Tank Farm Optimiser, substantial business and environmental benefits can be realised by wineries in their tank farm operations and end-to-end supply chain.

Bottling

The primary task of SolveIT Software's Bottling-Line Scheduler is to generate optimal production schedules for the wineries' bottling operations. The software uses advanced optimisation techniques for

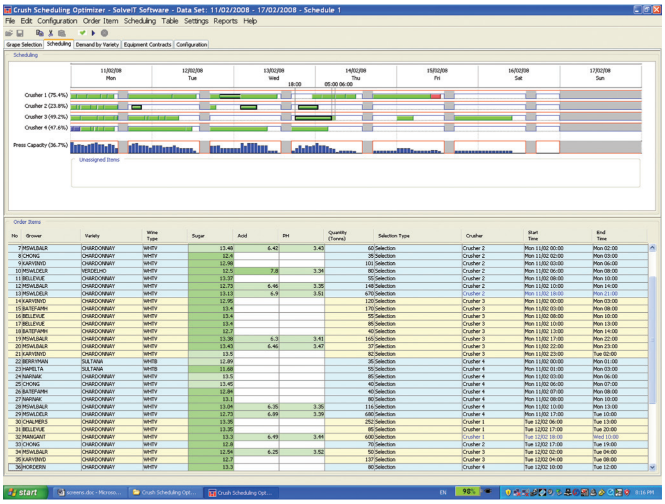

Figure 3. SolveIT Software's Crush Scheduler

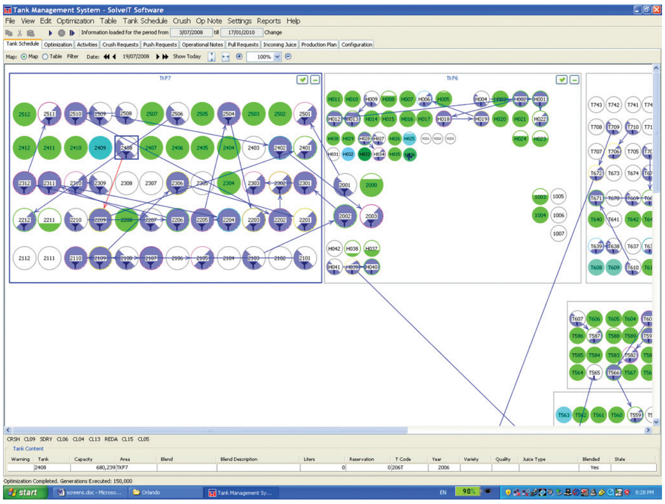

Figure 4. SolveIT Software's Tank Farm Optimiser

generating optimal production schedules. Opportunities for optimisation include manipulating the sequencing order, selecting which bottling lines to use, consolidating similar orders within the planning horizon, and suggesting changes to the requested dates that improve the overall schedule. Some of the key objectives are to maximise export and domestic service levels (i.e. DIFOT), maximising production efficiency, and minimising cost.

The user can see 'exception reports' where there are potential problems with work orders, or where there are items needing human attention. The user has full control over the optimisation process in that they are able to lock in manual decisions, set the business rules and constraints, re-optimise after making changes, and compare the current solution with an alternative plan.

The software also provides a *what-if* module that can be used to analyse strategic business decisions and events such as capital investment in new equipment, or to look at operational decisions like adding or removing extra shifts, or even for crisis management (what is the impact of a bottling line going down or key staff being ill). Reporting is provided on a number of levels and to suit different stakeholders; for example the daily bottling programs for execution, a report on particular wine blends, or a report on expected production efficiency.

The system generates robust schedules, such that small deviations from the planned schedule can be managed effectively (for instance, if excess volume is produced, or a machine breaks down for an hour, etc.). One aspect to robustness is to ensure sufficient lead time for supply of materials and for the release of wine from the tank farm. A significant part of the scheduling functionality is the ability for a user to adjust a schedule, rather than to completely re-schedule. This is important for two reasons:

- Schedules are generated on a daily basis. The previous schedule needs to be taken into account when looking to integrate new work orders. One of the major considerations is to minimise disruption to work orders close to being executed (i.e. schedule adherence), whilst allowing major re-ordering or the schedule for later orders to ensure the most optimal plan is produced.
- The user will regularly need to change the suggested schedule, to take into account user expertise and knowledge outside the domain of the software.

Where the user needs to change the schedule, a number of facilities is provided:

- locking an order in place;
- manually set the start date and production line for an order;
- setting Work Order specific rules (for example if a domestic Work Order must be produced by a certain date rather than being flexible); and
- updating the configuration of the production environment, including business rules and constraints (for example, to let the system know about a machine breakdown making a bottling line unavailable).

Once manual changes have been made, the schedule can be re-optimised, ensuring that an optimal schedule is still produced whilst taking into account changes to the production environment and human expertise.

Users are also able to generate an alternative to the current production schedule, with the system providing a comparison to help the user evaluate the impact of any changes. The comparison includes performance metrics (KPI's) such as any difference in the number of late orders, and changes to production efficiency and cost (this could include measures such as cost per unit, total production throughput, production line utilisation, etc). This allows the user to experiment with different schedules before committing to making any changes; for example, trying to incorporate a last minute export order without disrupting existing orders.

Environmental Factors

Further extensions to all of the presented modules (grape maturity, vintage planning, crush scheduling, tank farm optimisation, and bottling-line sequencing) include incorporation of water use and CO_2 emissions to allow wineries to better understand and optimise the trade-offs between climate change related issues, such as CO_2 emissions and water usage, and more traditional supply chain objectives like cost and volume. SolveIT Software's applications support both strategic and tactical decision making:

Strategic decision support includes:

- comparing different 'what-if' scenarios and their KPIs;
- understanding the trade-offs between competing KPIs, such as average production cost vs. water use vs. CO_2 output;
- simulate changes to plant equipment, warehouse capacities, product volumes and mix, etc.; and
- simulate changes to production and supply chain constraints and business rules.

Tactical decision support includes:

- planning and scheduling of production and supply chain activities;
- optimising plans and schedules, e.g. for minimum CO_2 emissions;
- managing unexpected events through dynamic re-planning and re-scheduling; and
- understanding why decisions are optimal, i.e. decision visibility and transparency.

These better-informed decisions address the effects of climate change through:

- reduced energy usage resulting in reduced CO_2 emissions;
- reduced water use; and
- increased ability to respond to adverse climate events, e.g. a heat wave or extreme rainfall.

SolveIT Software's capability to handle complex, multi-objective supply chain problems allows companies to optimise their operational activities within the context of their business rules and constraints, e.g. quality standards, cost control, customer delivery requirements, etc. For example, through the proper use of SolveIT Software's Tank Farm Optimiser, a winery can realise four significant benefits:

- Decrease in the number of wine transfers, which equates to:
  - a decrease in water consumption;
  - a decrease in carbon footprint;
  - an increase in product quality;
  - a decrease in product loss;
  - a decrease in maintenance spends; and
  - an increase in safety.
- Decrease in labour requirements.
- Increase in tank utilisation by reducing 'free working space' on the tank farm.

Each of these three benefits is covered in detailed in the following subsections.

Decrease in wine transfers

Most wineries use a white board approach for making wine transfer decisions with no forward visibility or optimisation. Tanks are assigned to winemakers, then each winemaker uses his/her own group of tanks (which are not optimal). According to our estimations approximately one third of the transfers are non-core 'pack up' and 'pack down' operations. By introducing a centralised Tank Farm Optimiser, the number and distance of transfers can be significantly reduced. SolveIT's Tank Farm Optimiser can simultaneously:

- reduce non-core transfers by 15% – 30% (for an overall reduction of 5% – 10% in tank farm transfers); and
- reduce the overall transfer distance.

A reduction of 5% – 10% in tank farm transfers for a tank farm utilising 100 million litres of water and performing 5,000 transfers/year equates to the following:

- A reduction of between 5 to 10 million litres in water consumption on the tank farm. This saving is achieved through fewer transfers, shorter transfers, and less 'water intensive' transfers.
- A reduction of between 5% – 10% in electricity consumption used for transfers, with corresponding reduction in carbon footprint.
- An increase in product quality (each transfer erodes product quality).
- A reduction of between 5% – 10% in wine loss. On average, 300 litres of wine are lost during each transfer. If we assume an average cost of wine of $1.5/litre then the annual financial benefit is between $112,500 to $225,000 ($450 × 250 to $450 × 500).
- A significant reduction in maintenance spend, as a reduction in transfers equates to a reduction in plant maintenance.
- An increase in work safety as each transfer carries a risk of injury.

Decrease in labour

Weekly payroll at a large winery may contain up to 10% of overtime wages. This may happen due to a lack of proactive work assignments. In such a situation, employees often have little to do Monday through Tuesday, and then too much Wednesday through Friday. Also, due to a lack of forward planning, considerable time might be spent setting up transfers (e.g. if two transfers use the same transfer path, it is more time-efficient and cost-effective to schedule them on different days). Otherwise, a lot of time is spent on finding an alternative path, which is often more expensive in labour setup time and water use).

Savings in this area can be achieved through 'labour balancing' (i.e. advanced planning of transfers) and through the reduction in changeovers from red to white in both transfer paths and destination tanks (which requires more labour and water for cleaning). Through forward planning of transfers and transfer paths, overtime can be cut by at least 50%, and fixed labour by at least 10%. For a tank farm with a $2 million annual payroll this equates into an annual financial benefit of $280,000 ($200,000 overtime × 50% + $1,800,000 fixed × 10%).

The realisation of this benefit would require process change in the handling of work orders at the winery. Specifically, the Tank Farm Optimiser would need to have the ability to delay or bring forward work orders within the week (i.e. move them plus/minus a few days) to achieve the 'labour balancing' effect.

Increase in tank utilisation

Tank farms keep a minimum of 'free working space', which may be anywhere between 10% and 13%. On a 100 million litre tank farm that equates to 10 to 13 million litres of empty space. Through the forward planning of wine transfers (rather than using the white board approach), this 'free working space' can be reduced by 1% to 1.5%., for a gain of 1.0 million to 1.5 million litres in capacity. This corresponds to an immediate once-off gain in new capacity (equivalent to several millions of dollars in capital investment) and represents one of the largest financial benefits of deploying the Tank Farm Optimiser. This increase in capacity will also provide the winery with a significant increase in throughput and decrease in average processing cost.

Conclusion

There are significant benefits in optimising the wine supply chain, which include the following:

- Moving away from spreadsheets, which usually grow in complexity as more elements are added.
- Provision of centralised applications that are maintained and supported.
- Provision of a scalable platform for future extensions.
- Straightforward integration with other applications for prediction and optimisation.
- Provision of integrated views (carrier, winery, etc.)
- Provision of integrated inputs (e.g. for Grower Liaison Officers and Logistics Coordinators).
- Provision of optimised capacity planning (i.e. automated 'smoothing').
- Automatic generation of robust, optimised production schedules that maximise service levels and utilisation, while minimising cost.
- Faster feedback to production planners, management, sales, and other interested parties for placing orders or requesting changes.
- Better use of available production capacity.
- Reduced risk of late deliveries due to production capacity issues, supply issues, or scheduling errors, and better visibility of potential risk.
- Schedulers can quantify the relative merits of different schedules, make informed decisions as to which scheduled to choose.
- Higher confidence in production schedules may allow running at lower inventory levels.
- Long-term planning from sales forecasts (e.g. assist with production capacity planning and production smoothing, supply planning for long lead-time items, and inventory planning, 'what-if' scenarios for strategic and operational planning, testing the impact of changes on business rules, infrastructure investment, overtime or extra shifts).
- Reporting on production (e.g. identification of capacity problems, identification of production or supply bottlenecks, high-level overview as well as information on specific orders).
- Reduction in overall transfers, leading to less water consumed in the wine production process, a smaller carbon footprint, less spillage, less plant maintenance, and increased safety.
- Reduction in labour requirements through labour balancing, a reduction in labour-intensive operations, and a reduction in overall transfers.
- Reduction in 'free working space' on the tank farm, leading to increased tank utilization, capacity, and throughput.
- Process improvement in the area of work order handling, by reducing paper handling and data duplication.

There are also a number of flow-on benefits, e.g. planners require less time to produce bottling plans, less chance of human error, identification of potential data problems, ability to handle dynamic changes to the schedule, whilst minimising the impact on existing Work Orders near their production date.

Acknowledgements

The authors would like to acknowledge the contributions of many individuals from Pernod Ricard Pacific, Fosters, Australian Vintage, and the Australian Wine Research Institute.

Reference

Michalewicz, Z.; Schmidt, M.; Michalewicz, M.; Chiriac, C. (2007) Adaptive Business Intelligence: Springer.