KIRONA™
Intelligent Mobile Solutions

A Beginners Guide to
Field Workforce Dynamic Scheduling

Contents

An introduction to job scheduling

If you are considering introducing an optimised and dynamic scheduling solution to allocate work to field staff more efficiently, then there are numerous things to consider. Factors such as the possible return on investment and strategic goals, the business process changes required to implement such a solution and the personnel required to implement and maintain the system, all need to be aligned.

The purpose of this documents it to give a broad introduction into optimised dynamic scheduling; assisting the reader with the understanding of how such a solution works. All optimised scheduling engines operate differently, using different algorithms to optimise job allocation. For illustration purposes, this document will occasionally focus on the specifics of Kirona's optimised scheduling system; Xmbrace DRS.

For the generic purposes of this document we refer to "jobs" and "field workers". Jobs can be any work carried out, outside of an office environment; a service, repair, inspection, interview etc. A field worker carries out the above jobs, working in sectors as diverse as healthcare, facilities management, engineering, business compliance etc.

Every company and every field based service is different

One of the most challenging aspects of mobile workforce management is getting the right job, to the right person, at the right time, with all the information and equipment they need to do that job.
The challenge is that within one organisation there are a host of "rules" that need to be applied in order to successfully allocate a job and there can be a variety of different work patterns for employees as well; shift patterns, skill sets/trades etc. For example, you could have all of the below requirements for scheduling in just one team:

- Individuals that work Mon-Friday 8am – 5pm, with 1 weekday off each fortnight and a Saturday shift 9-4pm that week instead
- Field workers that work on their own, but for certain job types must be sent the same work in order to work as a gang to fulfil that job (for operational or health and safety purposes)
- Jobs that need to be chained in a logical order and only become available to a field worker once the previous stage is complete e.g. a plasterer only being sent a job after the electrician finishes theirs, a district nurse will visit after an assessment by an occupational therapist etc.
- Jobs that have a specific start and end date with options for completion any time in that period, other jobs with specific SLA's e.g. 24 hr emergencies, other jobs that are appointments and must be completed within short windows of time e.g. 1 hr slots

Scheduling has multiple stages

From the creation of a job to the completion by a worker, and future jobs beyond that, there are numerous stages where the scheduling of work can be optimised to provide the best result to that organisation.

- If the customer requires an appointment, then the one offered to them should be offered based on the actual capacity of the workforce (not planned, or estimated), in order to ensure the appointment can be met and the customer isn't disappointed.
- If an appointment is created then the person booking it should know not only workforce capacity, but also the impact that appointment will have on job cost e.g. increase in universal travel time of operatives, in order to suggest an appointment that will minimise the cost to the service organisation.
- When the work schedule is laid out, it should have considered all the factors required to determine, which field workers can do that work i.e. skills, geography, equipment, and then allocate it to the worker that will cost the organisation the least e.g. minimal travel time to that and other jobs, allowing for more productive time
- The schedule should be monitored and jobs should be dynamically re-scheduled between staff continually to manage exceptions that take place e.g. jobs overrunning, adverse weather in order to ensure they are met on time (appointments, SLA's, end dates) and that costs are kept to a minimum i.e. if a job needs re-allocating from one worker to another then this should have minimal impact on travel costs

An optimised scheduling solution should take the above opportunities as well as any others that are presented, to ensure that the best result to that organisation. What is a best result is entirely configurable for that business division, as the scheduling system can be setup not just for the whole organisation but for various working teams that operate differently and have individual goals.

Under the bonnet of optimised scheduling

Unless a service organisation has a very small workforce or a very simple work process, then scheduling work will almost always take a significant amount of that organisations time, and crucially getting it wrong can have severely harmful effects on overheads and customer service. For many organisations these harmful effects are assumed to be a necessary and unavoidable consequence of their business. They are not.

Let's illustrate how difficult scheduling of work really is.

- If you have 5 field workers, carrying out 5 jobs a day at separate locations. Then there can be 3,125 different scenarios for allocating those jobs. If you have 100 workers carrying out 5 jobs, then that's 10,000,000,000(1billion) possible combinations of jobs, and this is a simple process!
- Often, allocating jobs isn't as simple as the number of workers and jobs; there are all sorts of other possible factors that determine it such as worker skills/trade, worker certification required (and the validity dates for the workers), worker languages spoken, job timeframe (appointment slit/site access), the region/geography, the worker's experience, parts, materials and equipment required and what the worker has in the van or depot, if the job needs multiple workers on site either one at a time or sequentially.
- Then there could be multi-dimensional objectives of the business – such as meeting the time sensitivities of the job (appointments/SLA) and reducing the costs of the jobs. Often these conflict and are made on judgement rather than empirical evidence
- Once a schedule has been created then that's just the start, the emerging day presents all sorts of challenges such as emergency jobs coming in, jobs overrunning, staff calling in sick; ultimately the perfect plan needs re-planning all over again multiple times that day ,with very little time available.

You are probably seeing just how complicated job scheduling is, if in fact you weren't already aware of it! Its complex, error-prone, expensive to manually manage, and often just doesn't get the best results; because there are too many factors to consider at any one point and nowhere near enough time to consider all options.

In summary, help is required in the form of an automated and dynamic scheduling engine to ensure the best results for the service operation in both cost reduction and improved customer service.

An optimised dynamic scheduler uses complex mathematical algorithms so that the process for the service team can be simple. The algorithms can carry out tens of thousands of calculations every couple of seconds to trillions within the space of an hour; to arrive at the best balance between customer service and operational cost, as defined by that service area. It presents in seconds what would take multiple schedulers hours or days, and limits manual intervention to an overview role.

The rest of this document will explore in more detail how a dynamic scheduling engine works and how it can be configured for the different work paradigms an organisation may need.

How an optimised and dynamic scheduling engine works

An optimised and dynamic scheduling engine is quite complex in how it works and how it can be configured, however at the basic level, the scheduling engine should do three fundamental things;

1. It should identify which field workers are suitable to undertake that given work
2. It should decide which field worker is the best person to carry out that work
3. It should refine this schedule to ensure that as time moves on the job is still done by the best person

How these processes work in detail is more complex, but explanations are offered below.

Process 1 - Which workers are suitable?

A piece of work, or job, often has a number of pre-requisites that must be met in order for it to be allocated to a field worker. These attributes are often the time constrain the work needs to be done by, the type of work and the geography.

These are then matched with attributes of workers to align the requirements of the work and the validity of the workers to decide who can do it. The number of factors can be virtually limitless, for example you could decide that geography, skills and time are important, but so are having the right equipment, speaking the right languages, having multiple staff available for one job. What these rules are is entirely configurable for and by

you; you could even set a requirement that this type of work can only be carried out by a field worker whose favourite colour is orange!

Typically the number of rules are best kept to under 10 as the more you have the less flexible you can be with the number of workers that can complete each type of job – leading to a smaller pool of possible resources and a greater chance of no resources being available or a greater cost because those that can spend more time travelling to the work, than onsite doing it.

Where there are multiple rules these are considered in sequence with a simple Yes/No. If all of the rules come out yes then this resource is suitable, if a single rule output is no then this worker is no longer considered. This is illustrated below in Figure 1.

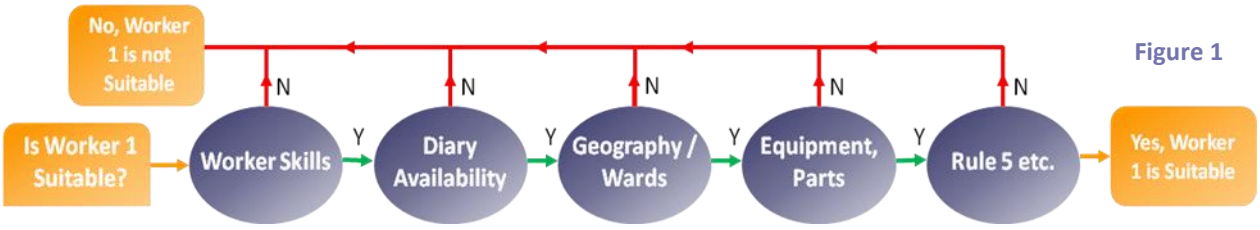

Figure 1

The above process happens multiple times with each worker and in this way there is an answer to whether a person can do the job.

Another factor to consider is that whilst most of these rules are mandatory, some can be optional. For example, you could have a field worker who has a primary and secondary skill. You want all jobs to be allocated to staff with the primary skill as that's the most efficient for your business, however if none are available, then the scheduler can allocate the work to someone with the secondary skill in order to act as cover. In this scenario there are 2 or more pools of labour suitable to do the job in a hierarchical order.

Output – The scheduler now knows all available field workers that are able to do this job.

Process 2 - Which worker is the best choice?

Once the full list of candidates that can fulfil that job are identified, the next step is to determine which one is the best choice; the optimum field worker.

As the rules in Process 1 have determined whether a worker can be onsite in the correct timeframe, it is assumed that customer satisfaction from a timekeeping perspective will be met. Now, with Process 2 we are optimising to ensure costs are now kept to a minimum. How costs are calculated to determine the most suitable task is a configuration exercise when the scheduler is setup to establish how an organisation wants it calculated.

The primary calculation in determining the optimum operative is typically the travel time. All travel is an overhead for an organisation which comes from a multitude of factors such as the fuel consumed, the depreciation of the vehicle used, the increased vehicle maintenance costs, mileage claims and the most significant; the opportunity cost of an operative. The opportunity cost simply means that if your field workers are driving then they aren't onsite doing the activity that you want them to do the most. Therefore reducing travel time reduces the vehicle based costs but also maximises their productive time. If an organisation can reduce operative travel by 10% for example, then they can in theory increase productivity by 10%.

To ensure that travel time is minimised the scheduling solution needs a travel matrix that understands the travel time from one job to the next. The crucial difference between a good scheduling engine and a poor engine is how this matrix works. A poor travel matrix will calculate distance as the crow flies; that is it will not take into consideration road distance and therefore can be easily fooled in calculating anticipated travel times.

A good travel matrix will have a very granular definition of travel with plots that are accurate to around 200m or less. This means it understands the true travel time, and is intelligent enough to know which road is the quickest route based on its length and the type of roads; similar to how a satellite navigation system calculates travel times.

See Figures 2 and 3 for an illustration of this. If a scheduling system uses a simple travel matrix then it is effectively working as shown in Figure 1. In this model the travel distance represented by the red arrow shows it is much more efficient to allocate Job 1 to Worker 1

Figure 2

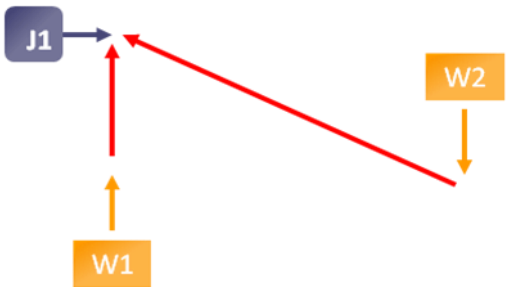

Figure 2 however is a very different story. This represents a good travel matrix that understands whilst W 1 is closest, there is a river in the way and the route to that is significantly longer than the route W2 would take; therefore the job should be allocated to W2.

Figure 3

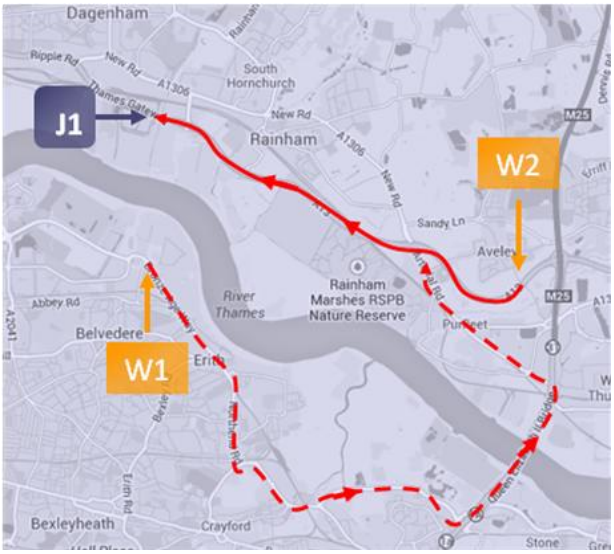

This is significant because a mistake in estimating the travel time means a worker may be unable to complete all their jobs in their given time frame leading to poor customer service and costly management of those jobs. It also means that jobs aren't being allocated with the shortest optimal travel routes, meaning field workers are spending more time travelling and given less time to carry out jobs.

Of course this problem becomes exacerbated the more jobs are added to the optimisation process. In Figure 3 below for example a simple matrix would have allocated J2 to W1 and J1 to W2. Actually because J1 will now be on the Way to J2 for W2, it makes sense to allocate both Jobs to W2 as there will be no increase in global travel time. This frees up W1 for other work.

Figure 4

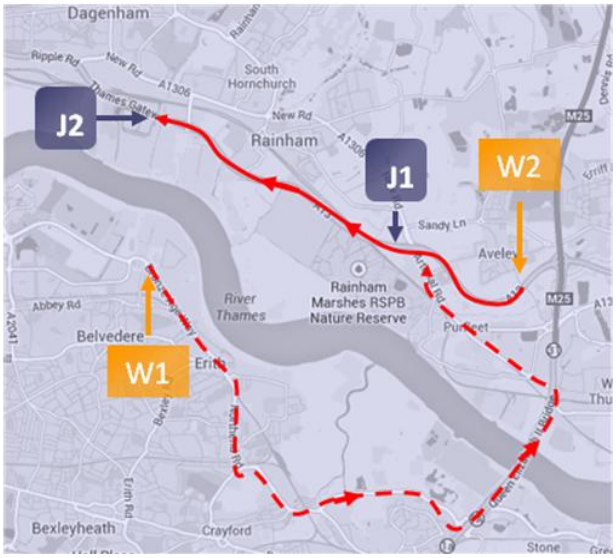

Whilst the above are simplified illustrations you can hopefully now see how the scheduler optimisation process is significantly improved by good street-level route optimisation. When optimising hundreds of jobs for tens of field workers over a period of days, the street-level travel matrix can make the difference between field workers having reduced their travel enough to carry out 1 extra job a day or not. Equally important is its ability to manage jobs over short timeframes such as a 1 hour appointment slot, as the scheduler knows how long the job takes and the travel time, giving an organisation a much better chance of meeting that appointment.

So far we've focused significantly on travel when calculating the cost. The scheduler however can also take into account of other costs when choosing work allocation such as the costs of the field worker against their colleague or a contractor, or if an field worker must be paid overtime to undertake a piece of work. These factors all contribute to minimising operational costs and maximising productivity through optimisation.

Output – The scheduler now allocates the job to the field worker that represents the lowest cost to the service.

Process 3 – Its changed, who's best now?

So far Process 1 and Process 2 have ensured the service schedule is optimised, by allocating the work to appropriate field workers with minimum cost. However, this plan isn't over until all the work is completed.

In the interim period between the creation of a plan and the plan being fulfilled, many factors can lead to this plan no longer being optimum. For example a plan may be created on Tuesday afternoon for the next day, however

first thing Wednesday morning emergency jobs are created that wreak havoc with this initial plan. These emergency jobs have to be allocated to the diary but the scheduler needs to find the best places for them to still maximise productivity.

An interesting concept here is that if a scheduling system merely tries to finds the best available space in a worker diary and adds that job to that space then this is unlikely to be the best solution as it doesn't consider the all of the jobs in this optimisation process. This is where the term dynamic scheduling comes into focus.

A dynamic job scheduling system constantly re-allocates work; it never has a static plan. If a new job comes in of higher priority then the diaries are re-scheduled to manage that job. If one worker over-runs with an appointment and will miss their next job, then a dynamic scheduler's rules re-allocate this job using Process 1 and 2.

Figure 6 below provides a hypothetical example where 3 workers have 7 jobs of varying start/finish times and durations, and varying travel times between them. Job 8 must be allocated to one of these available field workers. If all the other jobs remain stationary in these diaries then if Job 8 is allocated to:

- Worker 1 – the travel time for that worker will increase by 40 minutes
- Worker 2 – the travel time for that worker will increase by 60 minutes
- Worker 3 – the travel time for that worker will increase by 30 minutes

Figure 5

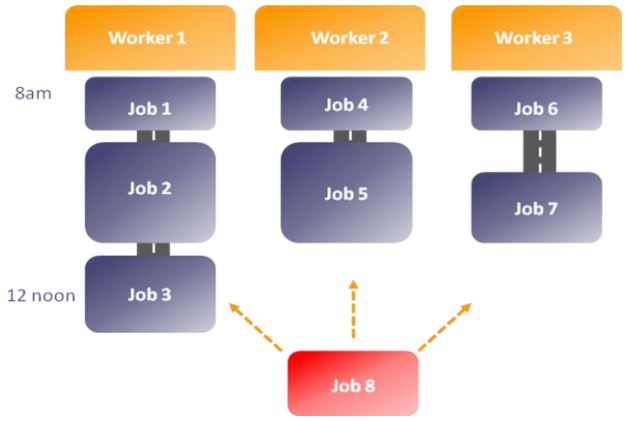

Logically you might think the job should therefore be scheduled to Worker 3 as that is the minimal increase in travel time but you can see in Figure 7 below it has gone to Worker 2.

This is because the scheduler hasn't just allocated that job onto an existing diary it has rearranged all jobs to find an even better travel time i.e. it has optimised all the jobs not just the new one, in order to find the lowest global cost. In the below scenario re-scheduling the jobs in this holistic fashion has increased the global travel time for all field workers by only 10 minutes, not 20 minutes.

Figure 6

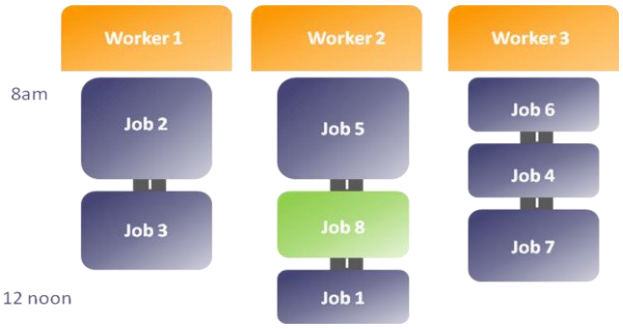

This is just a simple example. In reality a dynamic scheduling engine may carry out this optimisation for a host of reasons; a new job being created, a field worker starting/finishing a prior job late, a field worker being sick and unable to undertake any of their work, inclement weather etc. This point at which this optimisation takes place can also vary; it could be periodically, every time a new job is created be it emergency or planned, every time a job time overruns, every time a field worker explicitly notifies the scheduler via a mobile field worker application i.e. an inability to get access to a job, insufficient materials etc.

The larger the number of factors and workers the greater the calculation; this optimisation can number several trillion calculations (? x 1,000,000,000,000) in a couple of hours, achieving a result it may take a manual planner months or years!

Of course these factors vary from industry to industry. Typically organisations most impacted during an emerging day are those that are responsive services with lots of intra-day new work, especially if this is high priority. Services where jobs have a small time frame such as appointments are also greatly affected as being 10 minutes late can mean no-access to a site or the inability to do the job in the timeframe; jobs with wider timeframes are more forgiving.

Dynamic through Mobile Applications

A final point to mention about dynamic scheduling is the ability to carry out intra-day scheduling, and the success of intra-day dynamic scheduling is greatly improved if an organisation provides their field staff with a mobile application that allows them to be sent work in real-time.

By deploying a mobile application the scheduler can automatically send work to field staff without manual intervention. It can also use feedback from the mobile device to improve the dynamic optimisation. If an organisation provides field workers with mobile devices and only sends them one job at a time then this is typically the most efficient working practice for an organisation.

Figure7

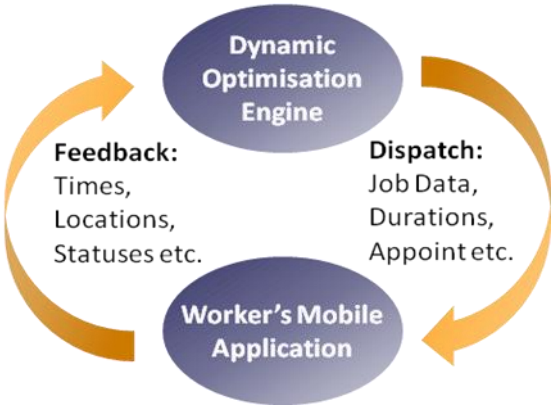

Updates about when the job is started / finished for example, trigger adjustments in the rest of the schedule before dispatching the next job to a worker. If a worker finishes a job earlier than expected then another job may be brought forward. This could be from a colleague's diary that day, or a planned job that was scheduled for some time later that month. Equally if a job over-runs and a worker's next job is in jeopardy, this could be allocated to a colleague. In this way the plan is constantly refined to manage the emerging day, and the results are more accurate as rather than making assumptions, the scheduler is being provided with verified and accurate data to make its decisions.

In this working process all the jobs in the plan remain fluid up until the point a worker requests or is pushed their next job at which point that job is fixed in their diary.

Output – The scheduler now may reallocate multiple jobs to the field workers that represent the lowest cost if factors change.

Appointment Booking

So far we've talked about core principles of dynamic scheduling. This section is going to focus on field services that require appointment booking, which whilst it used Process 1, 2 and 3, the time at which it uses them is very different.

Why appointed work is different is that it is an arrangement for two parties to meet at a specific time. For example if an engineer is visiting someone's home to carry out an emergency boiler repair, or an fitting satellite TV, or if a nurse is carrying out an outpatient visit then both parties must be available for that period of time. In effect this scheduling process is matching two diaries; one within the organisations control and one that isn't.

Most services are made when a customer rings to book a service via a contact centre. These calls have a small response time; typically 10-30 seconds and the shorter the better to reduce costs. In this timeframe it is very difficult for a call centre staff (CCS) to check the diaries of field workers to understand when someone is available and so the scheduling engine needs to present available appointment slots to the call centre staff in a different format.

Appointment slots can be anything that organisation is willing to work to e.g. 8:00-11:00, 11:00-13:00, 13:00-15:00 each day of the week, or different slots for different services on different days; maybe the above is for weekdays but on Saturday it's just 9:00-11:00 and 11:00-13:00. It is best for the CSS is this screen is launched directly from the CRM or other office system they are booking the appointment through to prevent having to logon to another system.

If a customer rings for a service type that has a 1 hour timeframe then CSS are made aware of which slots are available using Process 1, and a good scheduler will have optimised this to minimise costs using Process 2.

A very good optimised scheduler actually goes beyond this and carries out Process 3 for all the available appointment slots i.e. it tells the person booking the appointment not only the availability of field staff but also the increase in cost for booking each slot. This enables the CSS in just a couple of second to offer these "best" appointments with minimal cost for the service organisation, to the customer. If a scheduler doesn't do this and appointments are optimised after the booking may be inefficient for that time of day etc.

Figure 9 shows how this can be represented, and how it often is represented in our Xmbrace DRS scheduler. Red slots show a job cannot be completed and therefore cannot be scheduled in a in slot because of lack of resources, light green shows it can be completed and is within SLA, dark green is acceptable but outside of SLA. The stars represent the cost to the organisation with the most stars being the most efficient appointments. This therefore gives the CSS the information to understand what slots are available and which is minimal cost for the organisation to offer the customer.

Studies have shown that customers generally accept the first appointment offered to them 80% of the time, so this approach is very successful in reducing both call centre and field service costs, whilst maintaining the highest levels of customer service.

Figure 8

The final thing to note about appointment slots is the wider they are, the more efficient it is for an organisation, as they offer more flexibility for dynamically re-allocating work throughout the day. There is a temptation to think that the best customer service is also provided by offering a customer a narrow slot such as 15:00-17:00 but if that appointment cannot be made then it isn't, as you are more likely to let the customer down. Sometimes it's better with a broader appointment slot such as AM/PM. In the end good customer service comes down to a balance of offering a good service window but also keeping to that appointment.

Advanced scheduling

Up until now we've described the fundamentals of scheduling and you've probably got a good understanding of how it all works. But of course there are a variety of scenarios that increase this complexity. Some more advanced scheduling aspects are:

Cyclical and Recurring Work – A location may have a series of time dependant jobs that require spacing apart e.g. intervals of every month. The scheduler must consider all jobs in sequence to ensure adequate intervals.

Efficiencies – Some staff are more productive than others e.g. an apprentice may only be able to carry out 80% of the workload of an experienced worker. The scheduler should increase/decrease the standard minute value of a job proportionately to the worker experience to ensure their diaries are appropriately filled.

Forecasting – There is a desire to peaks and troughs in workloads and the capacity of the workforce to support this e.g. does a grounds keeping company have too few part-time staff in summer, or a heating organisation have too few boiler engineers in winter. The scheduler should present a forward plan both in live and in prediction based on historic trends.

Job Chains – Job A cannot be scheduled until Job B is done e.g. a decorator can't be sent to paint a wall until the electrician has carried out their job of rewiring it. The scheduler must find operatives who are able to do the jobs separately and then consider all of the combinations of workers to determine which is the best team for each time slot. Changes in the date allocated for A will influence B.

Manual Intervention – An organisation may want a level of intervention from the scheduler and allow planners to intervene and manually re-allocate work where necessary. The scheduler should allow a planner to manually change the job with new parameters or drag-and-drop the job into a new diary, locate a worker nearest to a job based on a map that is fed by their GPS position and allocate a job to them.

Multi-day work – This work takes more than one day to complete. The scheduler must consider the job and be able to either scheduler the work over course of several days and allow it to be "paused", or it should break the work into smaller jobs and schedule them.

Onsite Re-Booking – Being onsite often uncovers the full detail of work leading to that work being required to be re-scheduled or follow-on work to be created. In most cases the customer needs to be consulted and informed when the work will be done. The scheduler must be able to interact with a mobile application to enable field workers to book appointments in the field that are optimised, and confirm them with the customer.

Primary and Secondary Rules – The qualification process for allocating jobs to field workers should have optional rules in place for determining staff e.g. some field workers have primary and secondary skills, or work geographies. The organisation wants jobs sent to staff

with primary skills respected, but secondary skills can be used if none with primary skills are available to ensure a workload balance and increase scheduling tolerances. The scheduler should have an optional rule in Process 1 that weights the allocation to the primary factor and only allocated to the secondary if the primary can't be met.

Self Service – An organisation wants to reduce call centre traffic by enabling customer self-service appointments that still optimises scheduling. The scheduler appointment booking capability is embedded into the organisations website to present the customer with appointment options. These appointments can be presented one at a time to offer the most efficient for the organisation to the customer first.

Shifts – Workers may work on shifts that rotate. The scheduler should have individual diaries for workers that take their changing work hours into consideration.

Start Locations – Staff are allocated work from a start location e.g. a home or depot. But workers need to start at different location on different days. The scheduler should be able to multiple start and/or end location on separate days.

Teams – Two or more field workers who normally work independently may be required to work together on the same job e.g. someone working on a roof may require a second field worker onsite to foot a ladder a job. The scheduler must find operatives who are able to do the jobs separately and then consider all of the combinations of workers to determine which is the best team at what times

Quotas – An organisation may want to only fill field workers diaries to a certain percent of a certain job type e.g. there may be a high quantity of emergency work in a typical day and there is a desire to retain a space of a workers diary to ensure demand can be met. The scheduler should not allow more than a set % of these diaries to be used for 1 type of work.

Summary

Optimised and dynamic scheduling manages the entire lifecycle of work from its creation in an office system or contact centre to its completion by a field worker.

A good scheduling solution removes the subjectivity of a manual scheduler who doesn't have all the facts, with an optimised scheduler that has this information available and is able to make thousands of calculated decision in seconds.

It manages the multi-dimensional requirements of a work to ensure that the balance is fulfilled between operational efficiency and customer service. For example, if an organisation finds it has great customer service levels but high costs, then this balance can be quickly tweaked to align the service delivery to management goals overnight.

Building and maintaining a dynamic scheduling engine is no small feat. For a software company it requires years and millions in R&D to fulfil the criteria describes in this document. But for an organisation that implements it, the savings are tangible and highly visible; not only to the managers delivering the service on a day-to-day basis but also for senior management receiving reports on SLA compliance and reduced overhead

About Kirona

Founded in 2003, Kirona is recognised as a leader in the development and delivery of Field Force Automation software solutions.

We are successful because we tailor our approach to guarantee our customers achieve their business goals; increasing the number of jobs they carry out each day, reducing overheads such as fleet mileage and administration, and providing better customer satisfaction.

Over £20m has been spend in R&D of our product with over 25,000 field workers now relying on our software to undertake their work. Our portfolio of business applications enable them to streamline their service:

- Automatically schedule jobs and monitor their progress in real-time
- Reduce administration and travel in the mobile operative's day to provide more time for core tasks
- Locate and track your field based assets for operational and lone worker purposes

These applications can be fully integrated with their existing infrastructure, or be rolled-out as independent solutions, with the option of being hosted in Kirona's Private Cloud. In both scenarios the workflow is seam-less from the time a customer requirement for a service is created, to its delivery by the organisation's workforce.

Contact Us:

W: www.kirona.com
E: info@kirona.com
T: +44 (0)1625 585511

KIRONA™
Intelligent Mobile Solutions