

The A to Z of Mobile Workforce Scheduling Optimization

A ClickSoftware White Paper

November 2009

Copyright Notice

Copyright © 2009 ClickSoftware Technologies Ltd. All rights reserved.

No part of this publication may be copied without the express written permission of: ClickSoftware Technologies Ltd.

Publication Notice

The information contained herein does not constitute a warranty of any kind. ClickSoftware Technologies Ltd. reserves the right to revise this publication and make changes without notification. ClickSoftware Technologies Ltd. assumes no liability for losses incurred as a result of information contained herein.

Table of Contents

- The A to Z of Mobile Workforce Scheduling Optimization3
- Background - Scheduling optimization is much beyond naive schedule automation3
- Fundamental scheduling process of a work order3
- Scheduling customer appointments6
- Scheduling service requests for commercial customers6
- Ongoing Optimization7
- Daytime optimization vs night time optimization8
- Scheduling a batch of lower-priority work orders9
- Dispatching the work orders to the field technicians9
- Getting updates from the field technicians and maintaining the schedule accordingly10
- Scheduling processes that use advanced optimization11
- Summary12



The A to Z of Mobile Workforce Scheduling Optimization

For organizations that are considering a mobile workforce scheduling solution, there are a host of things to consider; cost versus benefit, change management concerns, personnel needed to implement the project and so on. What we are honing in on in this document are the finer points of an optimized scheduling solution. The goal is to arm the reader with enough information, and the right information, to choose the most optimal solution and to prepare for the implementation of that solution.

The Dynamic Nature of Service Delivery

The scheduling of field resources is an immensely varied and complex problem to solve. Within one company there can be many different types of field resources, work types, shifts and customers. For instance, within one organization you could have all of the following situations (and more):

- Individual field resources who work four ten hour shifts a week that must respond to all orders within three days and can be pulled away to work emergencies.
- Field resources that work together in crews, and are relocated frequently to remote areas, that work on orders weeks in duration.
- Contractor resources that are managed just like other internal resources, but can only do certain types of work.
- Orders for top tier customers that must be responded to within a certain period of time, or by a certain type of resource. Responding late could result in penalties paid to the customer.

Finding Opportunities for Service Optimization

There are many points along the service delivery process where the process itself can be optimized. A few of the optimization points in the scheduling process are:

- When the customer appointment is booked it should be offered based on true capacity, not planned or estimated capacity, to ensure that the field resources will be able meet the commitments.

- When the work is ready to be scheduled the scheduling solution should consider all available and qualified resources and choose the best one based on their skills, location, and more.
- The entire schedule should be monitored up until or even into the day of service, looking for opportunities to fill in excess capacity with low priority work, shift orders around for better routing, or react to changes in the schedule such as cancelled orders or rescheduled appointments.

Any optimized scheduling solution should take these variances and opportunities for optimization into consideration and provide a solution that is tailored to the different scenarios within an organization; rather than forcing the entire business to conform to one scheduling process. In the chapters that follow we delve deeper into how an optimized scheduling solution considers the variances within a service business and takes advantage of the optimization points to deliver service delivery improvement.

Background - Scheduling optimization is much beyond naive schedule automation

Every service organization with a moderate-to-large mobile workforce perceives the service operation as a core competence and spends lots of energy on scheduling the daily work of its field technicians. As a matter of fact, this is not surprising, since scheduling is a very complex process and, if not done well, can have a devastating effect on the organization. Let's review a few factors that demonstrate how complex the scheduling process is:

1. When an organization has two technicians who each do two work orders per day, the number of options for assigning the work orders to the technicians is 24. When three technicians each do three work orders per day the number of scheduling options jumps to around 360,000. When the organization has a few dozen technicians, each of whom performs six to eight work orders per day, the number of scheduling options is huge -and you can just imagine the number of scheduling options when the organization has a few thousand technicians. Even when the invalid scheduling options (e.g., attending work orders during non-working hours) are removed, the amount of valid combinations remains unrealistically large.

2. When adding to this process the multidimensional objectives of the business - such as keeping the service level as high as possible, keeping the operational costs as low as possible, and being compliant with the regulations - creating and maintaining the schedule becomes even more difficult.
3. Last but not least is the overall dynamic of the business. Customers call at different times throughout the day and ask to be served immediately. Other customers call to cancel the service they have just booked. Technicians report their daily progress, and some technicians call in sick. These are just a few of the additional factors to consider.

By now it is probably clear that service operation - and more specifically, scheduling management - is very complex, error-prone, and expensive when managed manually. In other words, it calls for automation. And optimization is the main means by which an organization can automate its scheduling operations.

Optimization consists of the appropriate software tools and the best practices that implement these tools to create and maintain the schedule with minimal human intervention. This document explains the essence of optimization. It uses daily scenarios of stereotypical service organizations in order to demonstrate the key elements that eventually lead to optimization.

Fundamental scheduling process of a work order

Before we dive in deep to understand the different elements of optimization and how they are being used to address the various daily scenarios, let's understand the fundamental scheduling processes. In order to simplify the scheduling process, we divide it into two steps:

- **Step One** - Identify the technicians who are candidates for doing a given work order.
- **Step Two** - Determine which technician is the best candidate to attend to the work order.

These processes are being used in almost all the optimization elements that we'll describe later in this document. The following further elaborates on each of these two steps.

Step One - Determine which technicians can be considered as candidates to do the work order

Scheduling a work order starts with validating compliancy with various constraints that the business dictates. For example, the work order dictates the skills set required in order to do the job and the scheduling process ensures that only technicians with these skills will be considered as candidates to do this work order. The constraints can also be considered as "rules" and a typical service operation has approximately 10 rules to consider. Examples include:

- a. Match the area - Only technicians who work in the area of the work order can do the job.
- b. Match vicinity from the home base - A technician can do the job if he lives within a certain distance from the job.
- c. Match the working hours - Only technicians who work during the working hours required by the work order can do the job.

The validation process yields a yes/no answer regarding every technician in question. Only the technicians who are positively validated by all the rules are considered as valid candidates to do the work order. For example, from a group of 100 technicians, the rules may narrow down the list to five technicians who can do the job.

Takeaway - Rules determine which technicians are candidates to do the work order.

Step Two - Determine which of the technicians is the best candidate to do the work order

The next stage is to evaluate which of the five candidates is the best option - in other words, who the optimal technician for this job is. In order to evaluate the options, the scheduling process uses soft constraints, sometimes considered as "business objectives." For example, if candidate 1 is currently located within a 10 minute drive from the work order in question whereas candidate 2 is located within a 30 minute drive, it will make more sense to schedule candidate 1 to attend to this work order - he is the optimal option. In a typical environment, the business dictates three to five objectives. When looking at the overall results of the business objectives, we will see that there is no right or wrong, good or bad option; it's simply a matter of choice to decide which of the options is optimal. The following table demonstrates a sample list of business objectives with their results.

Business Objectives	Technician No. 1	Technician No. 2	Technician No. 3
Travel time from the work order [minutes]	10	60	30
Overtime work if attending this work order [minutes]	30	0	15
Technician's cost per hour [\$]	50	75	35

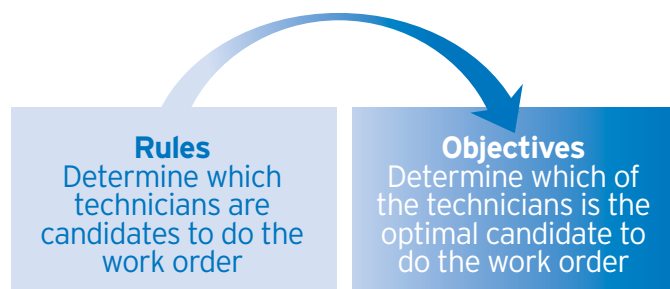
As the table shows, none of the technicians is preferred over his fellow workers in all three of the business objectives. The way to select the optimal option is by prioritizing the business objectives so that a more important objective will get a higher weight in the evaluation process. In addition, each business objective is associated with a monetary value that is used to evaluate the contribution (revenue minus cost) of each scheduling option. Eventually, the scheduling option that gets the highest score and contributes the most in terms of monetary value is the optimal option.

A different order of priorities (i.e., different weight and monetary value) will result in a different optimal solution and therefore the setup of business objectives is not complete until it goes through a balancing process that sets the right weights and monetary value to reflect the proper priorities. In reality, the balancing

process can be very tricky, as it relies on the work orders scheduled. The mutual effects of the business objectives on each other are very hard to predict. Experience shows that the most practical way to balance the business objectives consistently is to involve the business decision makers in the setup process and to run simulations of the real schedule. A comprehensive simulation process should probably consist of the following elements:

- Usage of historical records of real work orders.
- Mechanism that tests many schedules, each with a slightly different setup of business objectives.
- A way to articulate the results of each test so that a business representative can review a predefined set of key performance indicators, as well as several ways to visualize a snapshot of the schedule. After every iteration (simulation run), a business representative can more easily determine whether the scheduling results make sense or an additional iteration is required.
- A mechanism to allow changes to the prioritization of the business objectives. If the scheduling quality doesn't suffice, the business representative can change the order of priorities and run the simulation again.
- Once the business representative feels confident with the scheduling results, the setup can be recorded and is ready for production usage.

Takeaway - The business objectives determine which technician is the optimal candidate to do the work order and when that candidate should be assigned the work order.



Beyond the elementary scheduling processes, this document covers the following scheduling processes:

- Scheduling customer appointments
- Scheduling service requests of commercial customers
- Ongoing optimization

4. Dispatching the work orders to the field technicians
5. Getting updates from the field technicians and maintaining the schedule accordingly
6. Scheduling processes that use advanced optimization

Scheduling customer appointments

Suzan (the client) calls Better Serve (the cable TV provider) with a request to upgrade the cable TV package she uses. In order to fulfill Suzan's request, Better Serve will need to send a technician to install a new adapter at Suzan's home. At the call center, Jack (the customer service representative or CSR) takes the call. Suzan is ready to wait five to eight seconds for Jack to suggest a few time slots for the visit. During this time, the optimization system searches the scheduling board in order to find room for the new work order. It is extremely important to provide a reliable answer. Jack would like to refrain from overbooking the schedule, and thus every time slot he suggests to Suzan has to allow a successful schedule of the home visit.

It is possible to associate different types of time slots with different work order types or with different customers. For example, when Suzan calls, Jack identifies her and, according to her customer profile, offers her a time slot of two hours (for example 8:00-10:00, 10:00-12:00, 12:00-14:00 and 14:00-16:00). If Suzan calls and asks for a particular service, for instance to upgrade her cable TV package, the system associates a four-hour time slot for that type of service. In some cases, the regulator dictates certain time windows at which to start the work order (but not necessarily to finish it), and every customer who calls will get that type of appointment.

This process generally considers the scheduling board without moving any scheduled work orders, although some of the work orders allow room for a change in the schedule without breaking the commitment to the customer. As Suzan is going to select a time slot that is most convenient to her, the process should provide Jack with five to 10 different options that he can then offer her. In order to validate a time slot, the optimization process employs the match rules mentioned in Step One of the fundamental optimization process. As one can understand, the appointment booking process is limited by the time a typical customer like Suzan is ready to spend on the line with Jack. As a default, the process sorts the valid time slots in descending order, starting with the time slot that is closest to "now," i.e., the one that provides the service as soon as possible. If the business allows this process to last a few more seconds, it is possible to use the

business objectives (mentioned in Step Two of the fundamental optimization process) in order to sort the valid time slots in various options. For example, Better Serve can create a set of business objectives that focus on minimizing the operational cost so that the time slots will be sorted starting with the option that is cheapest to execute and ending with the most expensive option.

A typical CSR like Jack uses a customer-facing application such as CRM, an extension of ERP software or a homegrown application. These kinds of applications require integration with the scheduling application so Jack can run all the above processes while speaking with Suzan and using the customer-facing application.

Once Suzan selects her preferred time slot, Jack sends a message to book the time slot by scheduling a new work order. The schedule of the new work order restricts the availability of that time slot, and when the next customer calls, Jack will promise the same time slot only if there is enough time available for that customer's request.

Takeaway - Smart, real-time searching of the scheduling board is followed by immediate scheduling of the new work order.

Scheduling service requests for commercial customers

Jeff is the owner of Print Fast, a printing and document services shop. Print Fast has 10 different types of heavy-duty copiers. It is Tuesday morning and Jeff discovers that one of the copiers is broken. He calls Super Serve (a computer and office equipment service provider) and submits a report about the broken copier. At the call center, Andrew (the CSR) takes the call.

Andrew locates Jeff in the system, and according to the service contract, Super Serve must provide a technician on-site within eight working hours after receiving the call. In other words, there is a service-level agreement (SLA) clause of eight hours. Andrew initiates an automatic scheduling process that takes all the parameters into account, such as the location of Jeff's shop, the SLA, and the type of work, and employs the techniques of Step One and Step Two of the fundamental optimization process in order to schedule the new work order in real time. As Jeff waits, Andrew can afford about three to eight seconds until he needs to let Jeff know when the work has been scheduled and approve the schedule of the new work order. As long as the work order meets all the constraints (e.g., the SLA), Andrew simply needs to tell Jeff something along the lines of: "Your work order is confirmed and we'll be there on time."

During those three to eight seconds, an automatic process searches the scheduling board in order to find a space to schedule the new work order. The process can use a naïve approach and search in between the scheduled work orders without moving any one of them, in other words, search for the optimal "white space" for the new work order. This will obviously be the fastest way to schedule the new work order but the process may fail if the scheduling board is crowded, making it very hard to find space for another work order. A more sophisticated technique can slightly adjust the scheduled work orders within the limits of each, in order to increase the available space between them and thus increase the likelihood of successfully scheduling the new work order. For example if two work orders are scheduled for the same technician, one after the other, with a gap of one hour (excluding the travel time between the work orders), the system will not schedule a work order with a 90-minute duration, as the available gap between the two work orders is too small. If, however, the scheduling process "studies" the data of the two work orders, it can move each one of them slightly, so that the gap will suffice to cater to the new 90-minute work order whilst still meeting the commitments to other customers.

When it is necessary to "slightly adjust" the work orders, the schedule will be changed in one or more of the following ways: (1) By moving a work order to an earlier or later time; (2) by changing the arrangement of the work orders assigned to the same technician; or (3) by switching work orders between technicians. The scheduling process can also prioritize the new work order against the scheduled work orders and if the new work order has a higher priority, the scheduling process will be allowed to temporarily remove the lower-priority work order(s). That way the new space on the schedule is utilized by the higher-priority work order. For example, Jeff's call is considered a type of "break fix" work. If the work orders already scheduled are of a different type, such as periodic preventive maintenance, their priority could be lower (as the equipment can still operate). If the scheduling board is crowded, the scheduling process can remove one or more of the preventive maintenance work orders in order to schedule the break fix work order. Later, another automatic process will reschedule the work orders that have been dropped temporarily from the schedule. Both of these advanced techniques (moving a scheduled work order slightly and temporarily unscheduling a work order) are slower than the first technique, which sees the scheduling board as it is. Throughout

the implementation of such a solution, customers can balance their concerns about the slower speed of operation with the likelihood of completing the process with successful scheduling of the new work order.

Takeaway - Real-time scheduling of a new work order typically has a response time of a few seconds.

Ongoing Optimization

Background: The daily dynamics

The service requests of Suzan, Jeff, and hundreds of other customers populate the scheduling board and it becomes full of work orders. Speed is a critical factor when scheduling the work orders following Suzan's and Jeff's calls and being responsive sometimes results in a suboptimal schedule. As we are speaking about a dynamic environment, other changes affect the scheduling board from different perspectives. For example, some customers may cancel their booking, some of the technicians may call in sick, some of the work orders could start later than expected and some could be completed behind schedule. All these occurrences call for an automatic process that works in the background to improve the schedule and keep it running optimally at any point in time.


Background optimization: Definition

This background process works without any interaction from the schedulers and as a result it is sometimes called Background Optimization (BGO). BGO employs the scheduling techniques of Step One and Step Two of the fundamental optimization process, but with a more holistic approach. In other words, BGO considers all the scheduled and unscheduled work orders, as well as all the technicians of a specific territory or business line, and in a predefined frequency (e.g., every 30 minutes) attempts to improve the scheduling board.

For example, Better Serve as an organization is broken down into two geographical divisions: regions and districts. Every region consists of between 5 and 10 districts. In total, Better Serve has 1,500 technicians and region southwest consists of 120 technicians. If, on average, a technician performs six work orders per day, BGO considers a matrix of 720 work orders and 120 technicians to optimize for that region. The BGO can be set to optimize one region at a time, in a cyclical pattern, and thus region southwest will be optimized in a predefined frequency (e.g., every 30 minutes).

Opportunities to optimize: Rescheduling the unscheduled work orders

In addition to the scheduled work orders, BGO can also consider some or all of the unscheduled work orders of region southwest. As such, if one of the other scheduling processes unschedules a work order in order to schedule higher-priority work, the BGO can "pick up" the unscheduled work order and try to reschedule it. When a new gap has been created somewhere on the scheduling board (e.g., due to a customer cancellation), the BGO can identify it and utilize the gap for rescheduling.



Daytime optimization vs night time optimization

During the day, schedulers, technicians and automatic scheduling processes interact with the schedule and update it. If BGO considers all 120 technicians of region southwest, there is an increased potential for data conflicts between schedule changes that these users create and changes planned by the BGO. In order to avoid this problem and still run BGO during the day, the optimization divides the technicians into homogenous groups of 5 to 10 each, perhaps ensuring that each group consists of technicians who have similar skills. The BGO is then run on each of these groups separately, which reduces the likelihood of data conflicts. Dividing the technicians into groups is an automatic

process that works in conjunction with the BGO, therefore making it possible to run the BGO during the day.

On the other hand, the level of optimization is reduced as the matrix requiring optimization (number of work orders and number of technicians) becomes smaller. Service organizations that provide same-day service (the entire cycle of service from the time the customer calls until the technician finishes the work on-site occurs on the same day) find the daytime BGO useful, as it minimizes the impact of the daily changes and still helps to

constantly improve schedule quality. During the night, there are fewer interruptions (customer calls, updates from technicians, schedulers updating the schedule) and as a result the BGO can be set up to run on a more global basis with a longer run time, thus gaining better optimization quality. For example, the BGO can run on the entire region southwest, optimizing all work orders and technicians in the region. Some service organizations set up nighttime BGO to consider a scheduling horizon of a few days (e.g., 5), thus further optimizing the scheduling board.

Scheduling a batch of lower-priority work orders

Let's say Better Serve owns cable equipment that it maintains on an ongoing basis. According to its service policy, it gives a higher priority to work driven by customer requests, as these types of work generate money, whereas maintaining equipment is a means to sell more services to the customers. With that policy in mind, Better Serve decides that maintenance work should be done every quarter. It creates the work orders using one of the back-office systems and sends them as a batch through integration to the scheduling application. These work orders are available for the BGO to schedule them together with the other customer-driven work orders. According to the policy of Better Serve, it first schedules the customer-driven work and, around that, the maintenance work to fill in gaps of idle time. If the amount of maintenance work orders is relatively high compared to the amount of customer-driven work orders, it is possible to employ either or both of the following scheduling techniques:

a. Run a pre-scheduling process prior to running the BGO to identify a sub-list of maintenance work orders that can be scheduled. Alternatively, the BGO will need to scan a long list of work orders that will not be scheduled due to the limited capacity of the technicians. If we imagine that region southwest has 720 scheduled work orders per day and 5,000 maintenance work orders for the entire quarter, the BGO will work ineffectively if in each run it will need to review the scheduled work orders plus all 5,000 maintenance work orders. Alternatively, a pre-scheduling process will run once a day to scan the maintenance work orders marking about 100 of them and only these work orders will be considered for scheduling. As such, at the end of the quarter, Better Serve is more likely to complete all the planned maintenance work.

The pre-scheduling process uses its own algorithm, as selecting 100 work orders randomly would end up with a suboptimal schedule. Keeping in mind the business objective of using the maintenance work as "fillers" in between and around the customer-driven work, the algorithm can search for 100 work orders that are proximate to the locations of the technicians when doing the customer-driven work orders.

b. Run a periodic process (e.g., daily or weekly) that scans the scheduled and unscheduled maintenance work orders and increases their priority to reflect the intention to complete these work orders during the quarter parallel to the customer-driven work. Therefore, the priority of the maintenance work orders grows frequently until it is equal or even higher than the customer-driven work. In addition, the various scheduling processes will avoid removing the maintenance work orders in order to schedule more customer-driven work. This results in scheduling a chunk of maintenance work toward the end of the quarter.

Takeaway - Ongoing optimization during working hours accommodates the dynamic changes of the day and during non-working hours employs a global scheduling approach.

Dispatching the work orders to the field technicians

Sometime before Don (a Super Serve technician) starts a job, he should receive details of the work order through his mobile device. Better Serve is a unionized organization, and as a result the dispatching policy is to send details each afternoon of work scheduled for the following day. However, Super Serve has convinced their technicians that it's better for them to complete one work order before receiving the next work order (known as "drip-feed dispatching"). Common variants of drip-feed dispatching include:

- Sending the next work order X minutes (e.g., 60 minutes) before the technician needs to start traveling to the site of the work order.
- Sending the work orders for the next X minutes (e.g., 120 minutes) while allowing the technician basic visibility into the rest of his daily schedule, under the assumption that some of the work beyond X minutes may be rescheduled as the dynamics of the day dictate.

Drip-feed dispatching provides Super Serve with more flexibility to change the schedule up to the last minute when it dispatches the work order to Don (the technician). Full-day dispatching freezes the schedule a day in advance, but on the other hand, it provides technicians with visibility into the work for the following day.

Every service organization selects its preferred dispatching policy and, accordingly, an automatic process runs in the background and freezes the relevant work orders so that other processes can send them to the technician's mobile device. From that point on, the technician owns the work order and thus none of the optimization processes attempt to move it. When employing a hybrid drip-feed system, whereby the technician has basic visibility into the work orders beyond the next X minutes, the schedule of these work orders is not frozen and any rescheduling requires additional communication with the mobile device to allow the technician an updated view.

Takeaway - Once a work order is dispatched, its schedule is frozen and only the technician can update it.

Getting updates from the field technicians and maintaining the schedule accordingly

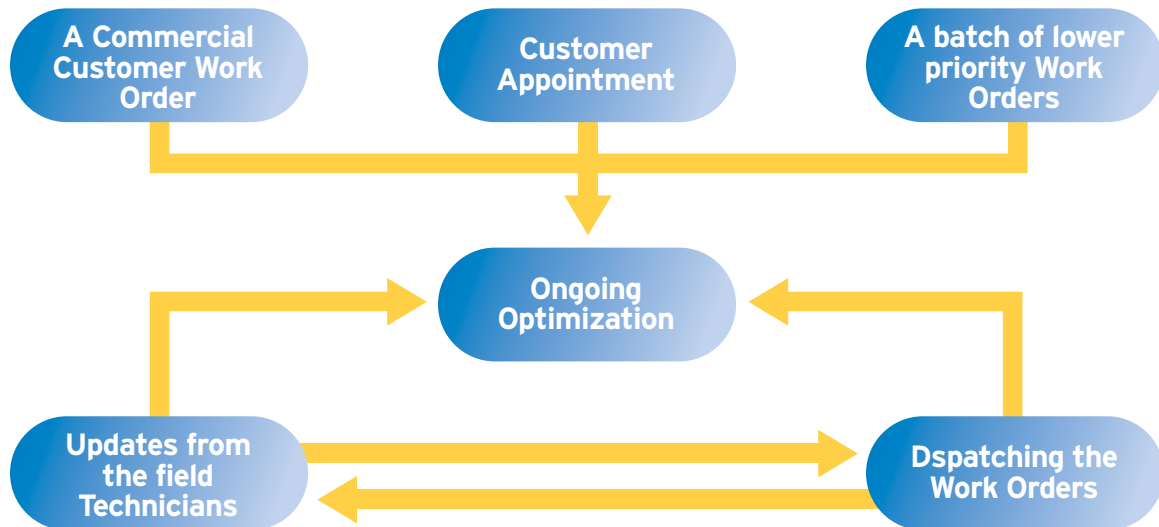
Don works as a technician for Super Serve. It is 11.30am on Wednesday and Don has just completed the work at Print Fast. The printer is up and running and Jeff is pleased with the level of service. As soon as he completes the work, Don uses his mobile device to report various details about the work, such as the parts he used, technical information regarding the copier and additional commercial information. From a scheduling perspective, Don updates the time he finished the work, and as soon as he synchronizes his mobile with the back office, the information about the work order updates the schedule in the following way:

- a. The original schedule includes information about the work order, such as planned start time, planned finish time and planned duration. Don's report updates the actual finish time and duration of the work order.
- c. If the actual finish time occurs before the planned finish time (that is, if Don completed the work faster than planned), then there is a gap in the schedule - white space that can be utilized to schedule another work order.
- d. If the actual finish time occurs after the planned finish time (that is, if Don completed the work later than planned), then his next work order is at risk of starting later than scheduled. This can create a domino effect on the rest of the day's schedule. To some extent, the risk may evolve to the point where Don won't be able to do one or more of the day's work orders.

- e. In the case of a late finish, an automatic scheduling process tries to "push" the rest of Don's work orders. It could be that the time constraints of each of these work orders allow this slight change without violating any of the obligations to customers.
- f. It could be that one or more of these work orders can't shift, as the move goes beyond the promised time slot. In this case, the work order can be marked as "at risk"; it will wait for the scheduler, who then needs to handle it by changing the schedule manually. Another option is to allow the same automatic process to temporarily remove the at-risk work order from the schedule and reschedule it, similar to the rescheduling of a new work order.
- g. In the case of an early finish, the process will work the same, but instead of pushing the rest of the work orders, the process "pulls" the work orders in order to utilize immediately the new white space in the schedule.
- h. Don's daily schedule is updated and the next work order that he receives on his mobile device will reflect the above process.
- i. The entire process of Don sending the update through his mobile device and waiting to receive his next work order should last about 30 seconds.
- j. If Don has idle time during the day, he can start a process that searches immediately for a work order that can be scheduled to him. The search will adhere to the service policy of each organization. For example, the system could search for a work order (scheduled or unscheduled) that is marked "in jeopardy" of missing its due date. If no work order has been found, then the system could search for a work order that is scheduled to Don later today or tomorrow. On a regular basis, the ongoing optimization and dispatching processes ensure that Don will have his next work order available on his mobile device at any point in time. In reality, sometimes there is a local failure in one or more of these processes. As a technician's idle time is very expensive, it makes sense to govern the automatic dispatching and allow Don himself to start such processes.

Organizations like Better Serve that dispatch the entire day's work in advance don't need to employ this automatic scheduling update process, as no scheduling process affects the same-day schedule.

Takeaway - Real-time reports from the field update the actual start and completion time of the work order and trigger adjustments of the rest of the schedule before dispatching the next work order.



Scheduling processes that use advanced optimization

So far we've demonstrated the elementary processes of optimization through examples that use a single technician, a single work order and basic daily routines. Using this as a basis, it is easier to understand the concepts behind more advanced aspects of optimization. This document doesn't provide detailed coverage of these aspects, but it offers insight into the complexity of handling advanced optimization. Some of the aspects of advanced optimization include:

- a. Crews - Two or more technicians may be allocated to work together for a dynamic period. The scheduling considers them all as crew when scheduling the work order, as well as considering the availability of each one of the technicians when he works for different crews or as an individual.
- b. Multiday work - A work order may be scheduled to span over multiple days. One of the challenges of scheduling this type of work order is calculating the duration of the work by considering elements such as the technician's availability on each day and the travel to and from his home base.
- c. Chains of work orders - Some work orders are dependent on other work orders. Examples of two types of such dependencies are:
 - i. Time dependency - Schedule work order A so that work order B will start only after work order A has been completed
 - ii. Resource dependency - Schedule work order A and work order B to the same technician/to two different technicians.
- d. Soft boundaries - Usually, technicians work within rigid territories. However, in some cases, it is more effective to schedule a technician from the adjacent territory, as he is closer than any of the technicians of the original territory.
- e. Relocation - Technicians may be temporarily moved from their default territories to a different territory to balance the workload between the territories.
- f. Home base vs. working area - A technician may live in location X, but most of his work is located in and around area Y. The work may be scheduled to take into consideration the route from X to Y by scheduling the technician to work orders that are located along his route.
- g. Lunch break - Most organizations allow their technicians to take a break for lunch. The timing of the break has some guidelines (e.g., not before 12:30 p.m.), as does the duration (e.g., 45 minutes). The optimization process determines when each technician will get his break by considering the work orders assigned to him and ensures that he is shown to be unavailable for work during the break.

- h. Long scheduling horizon - When the nature of the business is predominantly based on single work orders and single technicians, the scheduling horizon tends to be for a couple of days only (e.g., three to five days). When the work consists of multiday work and multistage work, the scheduling horizon can reach three to four weeks.
- i. "Flexible constraints" - In some cases, the scheduling policy doesn't fit the dynamic of the business and as a result many of the work orders cannot be scheduled automatically. For example, after a storm, there could be numerous power outages and if the volume of customer calls is very high, the predefined scheduling policy that was set to work on standard days will fail to handle all the work orders automatically. If the organization has a policy for extreme conditions, it can represent this policy in the system and initiate it whenever necessary.
- j. Capacity reservation - In some cases, the organization would like to ensure that resource availability is balanced between different types of work orders. For example, if an organization provides three types of service to its customers, it would like to ensure that resources will be available for all three services and avoid a situation where all the resources are utilized for one type of service, due to the daily distribution of customer calls. In other words, the organization is stipulating a given service level for every type of service.

Summary

Scheduling optimization deals with the entire life cycle of the work order, starting with the initial appointment requested by the customer or when the organization creates a maintenance work order. It ends with the successful completion of the work by a field technician. The dynamics of the service operation, the multidimensional business objectives and the amount of scheduling options call for automation of the scheduling processes.

A good scheduling optimization solution removes the personal subjectivity of the schedulers from the process, ensures consistent execution of the work, creates alignment with the management direction, sustains scalability of the business, and most importantly, increases the profitability of the organization. It takes years of experience and constant investment in R&D to build a world-class optimization system that can do everything described in this document. It is sometimes hard to see how each process works "under the hood," but fortunately the results of optimization are tangible and one can clearly see the resulting schedule and measure its quality.

About ClickSoftware

ClickSoftware (NasdaqGS: CKSW) is the leading provider of workforce management and service optimization solutions that create business value for service operations through higher levels of productivity, customer satisfaction and cost effectiveness. Combining educational, implementation and support services with best practices and its industry-leading solutions, ClickSoftware drives service decision making across all levels of the organization. From proactive customer demand forecasting and capacity planning to real-time decision making, incorporating scheduling, mobility and location-based services, ClickSoftware helps service organizations get the most out of their resources. With over 150 customers across a variety of industries and geographies, and strong partnerships with leading platform and system integration partners, ClickSoftware is uniquely positioned to deliver superb business performance to any organization. The company is headquartered in Burlington, Mass. and Israel, with offices in Europe, and Asia Pacific.

Contact Information:

For more information about ClickSoftware
visit www.clicksoftware.com
or email sales@clicksoftware.com

Offices

NORTH AMERICA
ClickSoftware (Corporate)
35 Corporate Drive
Suite 140
Burlington, MA 01803
Tel (888) 438 3308, (781) 272 5903
Fax (781) 272 6409

EUROPE
ClickSoftware
The Priory
Stomp Road
Burnham, Bucks
SL1 7LW
Tel + 44 (0)1628 607000
Fax + 44 (0)1628 607001

ClickSoftware Central
Europe GmbH
Hanauer Landstr. 293 B
60314 Frankfurt am Main
Deutschland
Tel + 49 (0) 69 489813 0
Fax + 49 (0) 69 489813 99

ASIA-PACIFIC
ClickSoftware Technologies Ltd
Azorim Park, Oren Building
94 Em-Hamoshavot Road
PO Box 3697
Petach-Tikva 49527
Israel
Tel + 972 3 765 9400
Fax + 972 3 765 9401

ClickSoftware Australia PTY Ltd
Level 1
256 Queen Street
Melbourne, Vic 3000
Tel +61 (0)3 9946 6400
Fax +61 (0)3 9946 6401

ClickSoftware Japan K.K.
15F Cerulean Tower, 26-1
Sakuragaoka-cho, Shibuya-ku
Tokyo, Japan
Tel +81 3 5456 5661
Fax +81 3 5456 5511

ClickSoftware India Pvt
Vatika Atrium
3rd Floor, Block-B
Sector - 53, Golf Course Road
Gurgaon - 122 002
Haryana, India
Tel +91 124 4311 188