

Operations Practice

Smarter, faster asset allocation for more coverage at lower cost

With careful planning, the latest technologies mean businesses can grow their asset bases without growing their cost bases.

by Guy Benjamin, Tod Camara, Joseph Faucher, Mayuri Singh, and Kumar Vivek



It's a common phenomenon across industrial companies: leaders looking for opportunities to slash the operational costs of managing physical assets, even as they make bold moves to install yet more assets. For example, renewable-energy companies have increased wind-turbine installations across the U.S., while telecom companies around the world are building out new 5G networks. As these efforts expand, the need to optimize and reduce operational costs is only going to intensify.

As many leaders already know, one of the best ways to keep costs in check when installing more physical assets is to increase the number of assets covered by their field-force workers—operators, engineers, technicians, and the like. After all, labor typically represents the largest share of asset-management costs. Yet few businesses have really cracked the code for field-force allocation: some still just rely on local managers to assign people to assets, while others use more sophisticated methods, but often with less impact than they had initially hoped. What these methods have in common is that almost inevitably, they sacrifice long-term effectiveness in favor of easy, fast implementation.

For years, leading companies have wanted to apply cutting-edge advanced-analytics and modeling techniques, in an attempt to avoid compromises between quality and speed when allocating assets to the field. Until recently, however, an analytics-based methodology was too demanding to implement. But the latest advances in technology have made it possible to simplify and expedite the approach, thereby unlocking the potential of smarter, faster asset allocation.

Today's common approaches compromise impact

We evaluate asset-optimization methods along two dimensions: effectiveness (the accuracy and efficiency of optimization) and ease of implementation (the complexity of the approach

and time to implement). Each of the four approaches typically used today suffers from a tradeoff between these two variables (Exhibit 1).

Manual and local. Field managers use local knowledge to assign assets to their field-force workers to minimize drive times, maximize specialized expertise, and balance workloads across the workforce. While this method is extremely simple and generally easy to implement, it places a heavy burden on field managers, and ignores available data that could yield a truly optimal allocation.

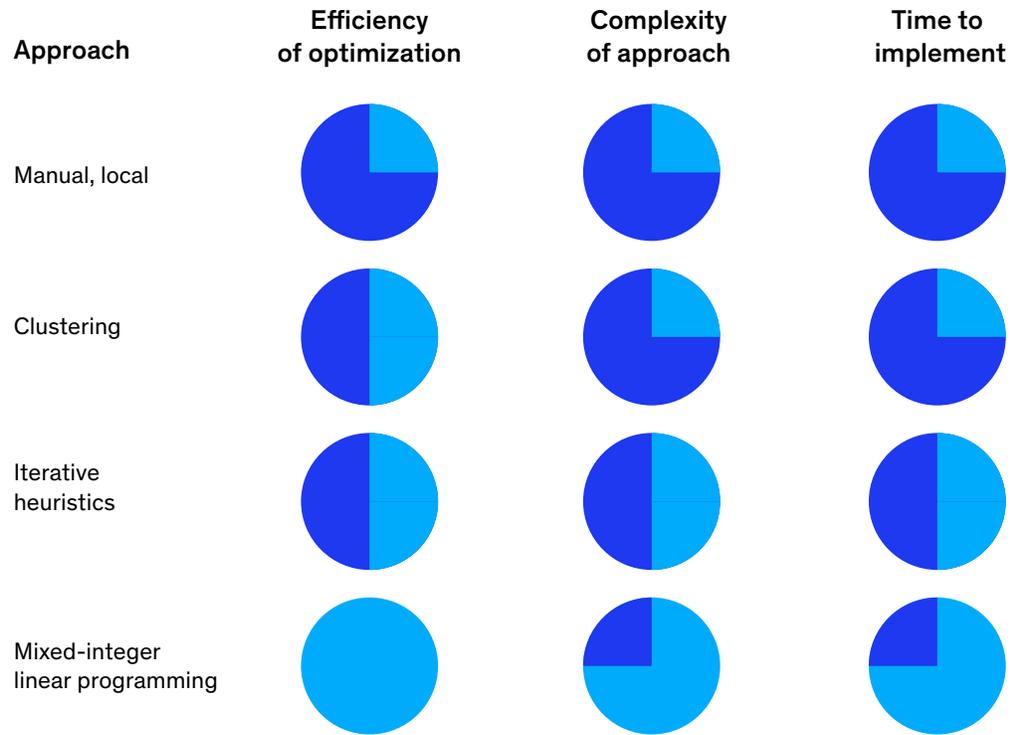
At a regional energy company, for instance, local managers typically assign field-coverage areas on the basis of industry benchmarks, such as the average number of tickets resolved by field workers. Managers spend only one to two days each year assigning coverage areas—a relatively small investment. But the rewards, such as they are, are equally small, producing suboptimal coverage areas and excessive drive times.

Clustering. Assets are grouped based on physical proximity, then assigned to the field force. Asset-workload and drive-time data can be incorporated via an iterative clustering process, allowing for asset allocations that better balance workloads and reduce drive times.

Yet clustering's impact often disappoints. At a large telecom company, this simple, analytics-backed approach was implemented in just six weeks, and relieved field managers of the burden of assigning sites. But it was only moderately effective in reducing the driving distance between the field force and work tasks, with drive times falling an average of only 10 percent. This was because, in general, clustering does not consider field workers' starting locations, making it impossible to actively minimize drive times. And because the telecom company's version of clustering did not optimally assign tasks across the entire field force, both drive times and staffing levels remained suboptimal.

Exhibit 1

Current optimization approaches are either inefficient or complex and time-consuming.



Iterative heuristics. Assets are assigned iteratively to the field force according to a tailor-made algorithm. This method incorporates the starting locations of the field force, and provides considerable flexibility to incorporate business constraints. However, as illustrated by a dispatch company, the iterative nature of the algorithm often results in suboptimal asset assignments.

Typical of iterative heuristics, the dispatch company started with a simple algorithm that assigned assets to the field force based on proximity, and then removed the assigned assets from further consideration. As the process continued, the remaining unassigned assets were interspersed among existing clusters.

The output seemingly met the company needs, requiring eight weeks of work to achieve a 20 percent reduction in average drive time. But it yielded diffuse, overlapping coverage areas that required more drive time for field-force workers assigned later in the process: the last 10 percent of the field force to receive assignments were given longer routes that had not yet been assigned. As a result, their drive times were significantly higher.

Mixed-integer linear programming. Assets are assigned to the field force using a mathematical formulation based on asset workloads, business-specific constraints, and drive times. Although mixed-integer linear programming (MILP) is the most complex and time-consuming of the asset-

optimization approaches, it provides the most flexibility by incorporating business constraints. It also yields the most optimal allocation across all assets by considering the entire solution space at once—no sites or field force are removed from consideration during the optimization process. Moreover, once the time-consuming step of developing the complex mathematical formulation is complete, the approach can be rapidly deployed.

A South American energy company applied MILP in redesigning its coverage areas to minimize field driving times and associated labor costs. While the effort reduced drive times by 30 percent, it also required 300 development hours across an intensive 12-week effort.

A simpler, more advanced MILP for faster, easier optimization

Simplifying the underlying MILP mathematical formulation can greatly expedite the method while maintaining its integrated view of the problem to be optimized. How? An adaptable, mathematical formulation that focuses on reducing drive time eliminates the time-consuming setup for MILP optimization. It also allows organizations to quickly and easily incorporate new business constraints, enabling them to tailor the approach to their needs—such as incorporating fixed limits for travel distance or coverage radius.

Moreover, this advanced method can incorporate a more detailed, realistic understanding of how the field force spends its time, adding administrative task time, maintenance workload, and additional job responsibilities to the usual drive-time data. At most large businesses, the inputs come from readily available sources:

- Asset and field-force location data, such as addresses or latitude-longitude coordinates
- Historic site-level workload data as a proxy for future expected work
- Field-force salaries

This advanced method is potentially useful to a wide variety of businesses that rely on their own field workers to maintain dispersed physical assets. Businesses ranging from telecoms and energy utilities to railroads, airlines, waterways, and even elevator and escalator manufacturers all face the problem of allocating workers to assets requiring service.

A large North American conglomerate shows the impact from this newer technique. Traditionally, the company's managers had allocated assets to their field technicians manually, without leveraging any analytics-backed methodologies. As a result, the average number of assets maintained by each technician is significantly below industry benchmarks, leaving the current workforce unable to support the planned expansion in the company's asset base.

By replacing manual allocation with advanced MILP, the company can reduce drive times by 20 percent, and trip distances by 25 percent (Exhibit 2). More importantly, these improvements translate into sufficiently large capacity increases that the company will be able to absorb the entire asset expansion using its current workforce (Exhibit 3).

How companies should build an advanced asset-management tool

Optimizing asset allocation using advanced, simplified MILP involves four steps, illustrated by the example of a telecom company (Exhibit 4).

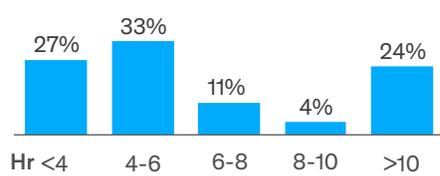
Step 1: Prepare data. As is typical of many companies with large asset networks, the data essential for the telecom company's MILP approach was distributed across several systems that did not communicate well with one another. The company's first step was therefore to identify the necessary data related to asset location, field-force location, and workload across its systems. It then created a key to map these data into a common data source. Having a single data source enabled the company to

Exhibit 2

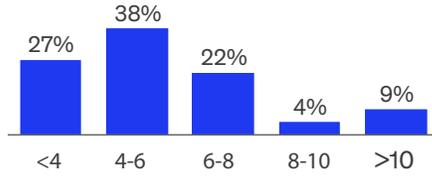
Advanced mixed-integer linear programming (MILP) reduces drive times and trip distances.

Weekly drive time

Current state
% Technicians



Future state



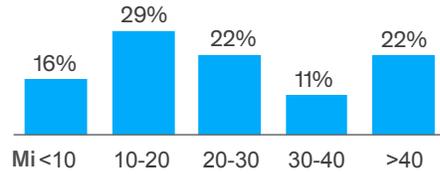
Impact



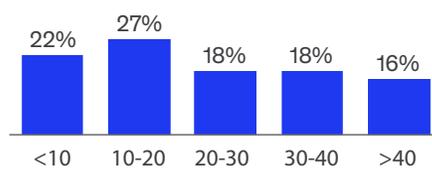
20%
drive time
reduction

Average trip distance

Current state
% Technicians



Future state



Impact



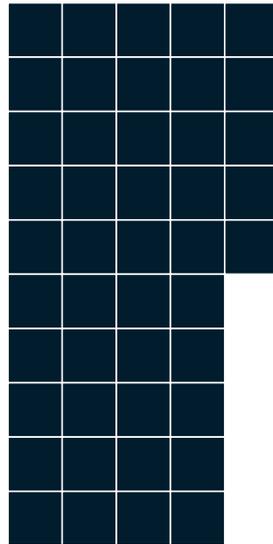
25%
trip distance
reduction

Exhibit 3

Reduced drive times and trip distances mean each worker can support more sites.

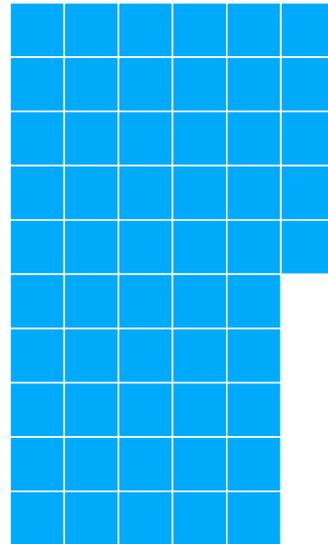
Current state:

~45 sites per technician

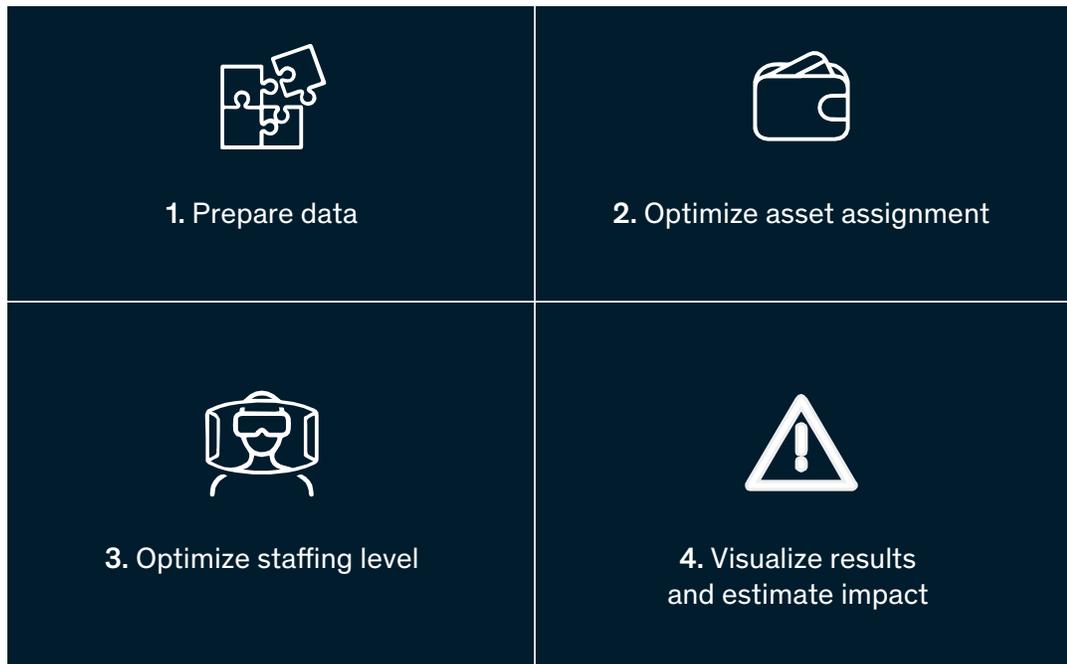


Future state:

~55 sites per technician



Take four steps to build an advanced asset-management optimization tool.



create a matrix, mapping the driving time to each asset for each field technician.

Step 2: Optimize asset assignment. The next task was to customize the MILP algorithm to the company's use case. The main challenge at this stage is translating complex business objectives (such as evenly distributing workloads or coverage areas) into a simple mathematical formulation. The telcom company opted for a basic mathematical formulation with the aim of reducing drive times. The company then incorporated its business objectives as constraints on the optimization algorithm: for example, it wanted to ensure that all field-force workloads would fall within a desired range. By running this optimization using the drive-time matrix and the workload data collected by the company, managers could reassign assets to achieve minimum drive times across the entire field force.

Step 3: Optimize staffing level. The third step adds a workforce component to ensure that staffing levels are just sufficient to support the asset network, without overstaffing any location. Creating this optimization provided deeper analysis of the geographic breakdown of the company's current workload, and how well the field force had been allocated to complete that work. Over time, the company can apply these insights to decide how to shift the size and placement of the field force to better align its asset footprint with its planned growth.

Step 4: Visualize results and estimate impact. The final step turns the output into an easy-to-use format so that leaders and managers can make improvements and respond more quickly to changes. The telecom company used interactive geoplots to map asset assignments, giving

managers a detailed view of their coverage areas. By reviewing these areas with their field technicians, managers were able to better support the field force in adapting to new assignments so that drive-time reductions could be achieved more rapidly. On an ongoing basis, the new visualization of workload across the field force enables the company to plan future initiatives more efficiently by better leveraging employee availability.

Companies no longer need to make the tradeoff between quality and speed when allocating assets to their field force. The new asset-optimization solution provides the best of both worlds through simple, adaptable mathematical formulation and nevertheless incorporates practical business constraints. With asset allocation that is both smarter and faster, companies can aggressively roll out and support new physical assets while maintaining low operating costs.

Guy Benjamin is an associate partner in McKinsey's New York office, **Tod Camara** is an engagement manager in the Chicago office, **Joseph Faucher** is an associate in the Stamford office, and **Mayuri Singh** is a senior knowledge analyst in the Gurugram office, where **Kumar Vivek** is an analytics specialist.

Copyright © 2019 McKinsey & Company. All rights reserved.