Solving the On-Demand Challenge:

How a Route Optimization solution can help you manage your logistics

It doesn't matter if you transport people, deliver food, or pick up and drop off courier packages, you need a smart solution to manage the logistics of your on-demand fleet.

Routific

On-demand is everywhere

In a world of on-demand everything, it's never been more important to make sure your business is ready to deliver.

The on-demand economy presents some serious logistical challenges – challenges that an operations manager or even a roomful of dispatchers cannot solve efficiently. In today's market, logistics businesses are expected to be flexible enough to constantly react and adjust to real-time information. Whether it is the placement of a new rush order, last-minute cancellations, or other unforeseen events, your operations team must deal with these as efficiently as possible, in order to stay competitive.

Manual dispatchers are a thing of the past. Modern businesses that are ready to scale and disrupt the market are already turning to dispatch algorithms and routing software for help. But not all algorithms are created the same.

It doesn't matter if you transport people, deliver food, or pick up and drop off courier packages – you need a smart solution to manage the logistics of your on-demand fleet.[1]

This whitepaper explores six different algorithms that can be used to handle the logistics in a real-time dispatching environment. You'll see that a route optimization performs pretty darn well. So well, in fact, that it far outperforms many of the other algorithms out there. It has the potential to help businesses make more deliveries in less time, and improve profitability by 41%.

[1] We are assuming true "on-demand" companies that promise same-day pickup and delivery. Some on-demand companies have the luxury to offer next-day or scheduled deliveries. However subtle the difference, the former can be solved with dispatching algorithms described in this article, whereas the latter can only be solved with route optimization algorithms.
For a further treatise on this topic, and why companies that schedule pickups and deliveries in advance will be the true winners in this space, read: The Case Against Everything on Demand: www.recode.net/2014/09/15/the-case-against-everything-on-demand/)

Dispatching in the real world

Before there was on-demand food, laundry, dog-walking, massages, gasoline, or cannabis, the only players in the on-demand space were traditional courier and taxi companies.

Most of these businesses are so inefficient, it's hard to believe. Traditional companies typically employ a roomful of human dispatchers, each committing about 5 to 10 drivers to memory. Dispatchers then communicate with their drivers using ancient radio technology.

Dispatchers are often former drivers who have spent a lifetime on the road, because without modern software, knowledge of the streets is still a necessity. During our industry research, we were shocked to discover that many don't even use digital maps.

"How do you keep track of everything? The roads, your drivers, and the changing schedules?" we asked them.

The "best" dispatchers would tap their heads and say: "It's all in here, buddy."

The status quo in this industry is abysmal.

The good news is that this is changing quickly with the rise of disruptive on-demand startups. For a modern, on-demand company to be truly scalable, it needs to rely heavily on technology and dispatching algorithms. In most cases, even simple heuristics will get the job done, but the real question is: what margins are you leaving on the table?

As for traditional companies that still employ manual dispatchers, it is a matter of survival. Margins are bleeding. How do you stay competitive?

The Same-day Pickup and Delivery Problem

Locations 7

Packages 320 $50

Vehicles 5 4

Maximum score $16.000

To study the performance of different algorithms and heuristics, we ran them through a simulation model and compared the results.

In this simulation model, there are 7 locations. Packages have a pickup and dropoff location, as well as a "latest time of arrival" constraint. During the simulation, pickups are requested randomly at any time, and some requests may even get canceled if you don't pick them up quickly enough. You have 5 vehicles in your fleet, each with a maximum capacity of 4 packages.

The simulator awards $50 for each package that is successfully dropped off on time. The objective is to maximize your profits by delivering as many packages as possible using a centralized dispatching algorithm. In total, there are 320 packages spawned during the simulation, so the upper bound is $16,000.

The following are 6 high-level descriptions of popular dispatch algorithms. [2]

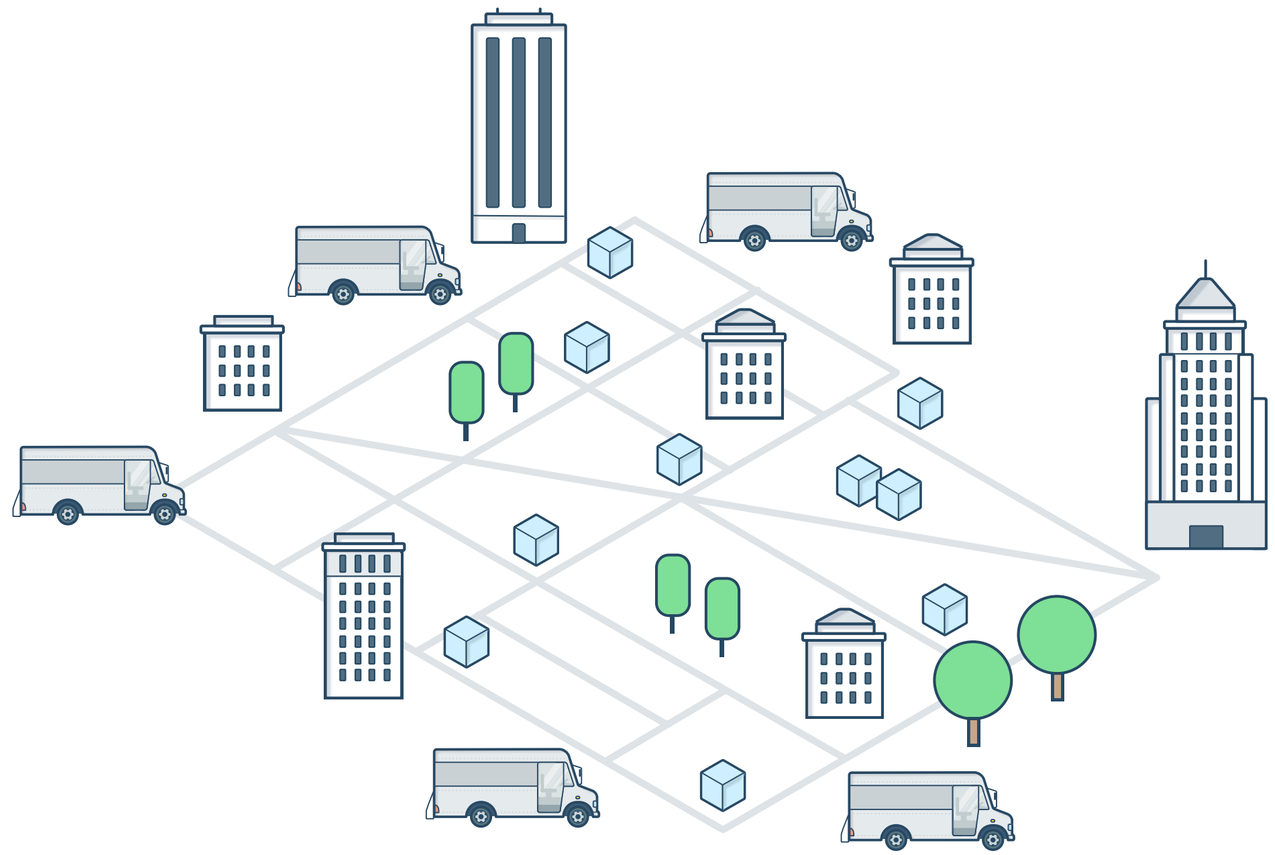

[2] You may argue that these are but simple heuristics, rather than "real" algorithms (except for the route optimization approach). Nonetheless, most on-demand companies - even some of the largest ones - simply resort to very basic greedy heuristics. However inefficient this approach may be, it gets the job done. While getting the job done with heuristics may have been a great leap forward compared to manual dispatching, in this increasingly competitive space, it becomes more and more important to pay attention to squeezing out efficiencies.

1. Pure greedy

Method:

The easiest method of dispatching your vehicles is to simply assign the closest available car. This is typically what the dispatchers at most taxi companies do. We'd venture to say that this is probably the most common approach for on-demand companies because of its simplicity and ease of execution.

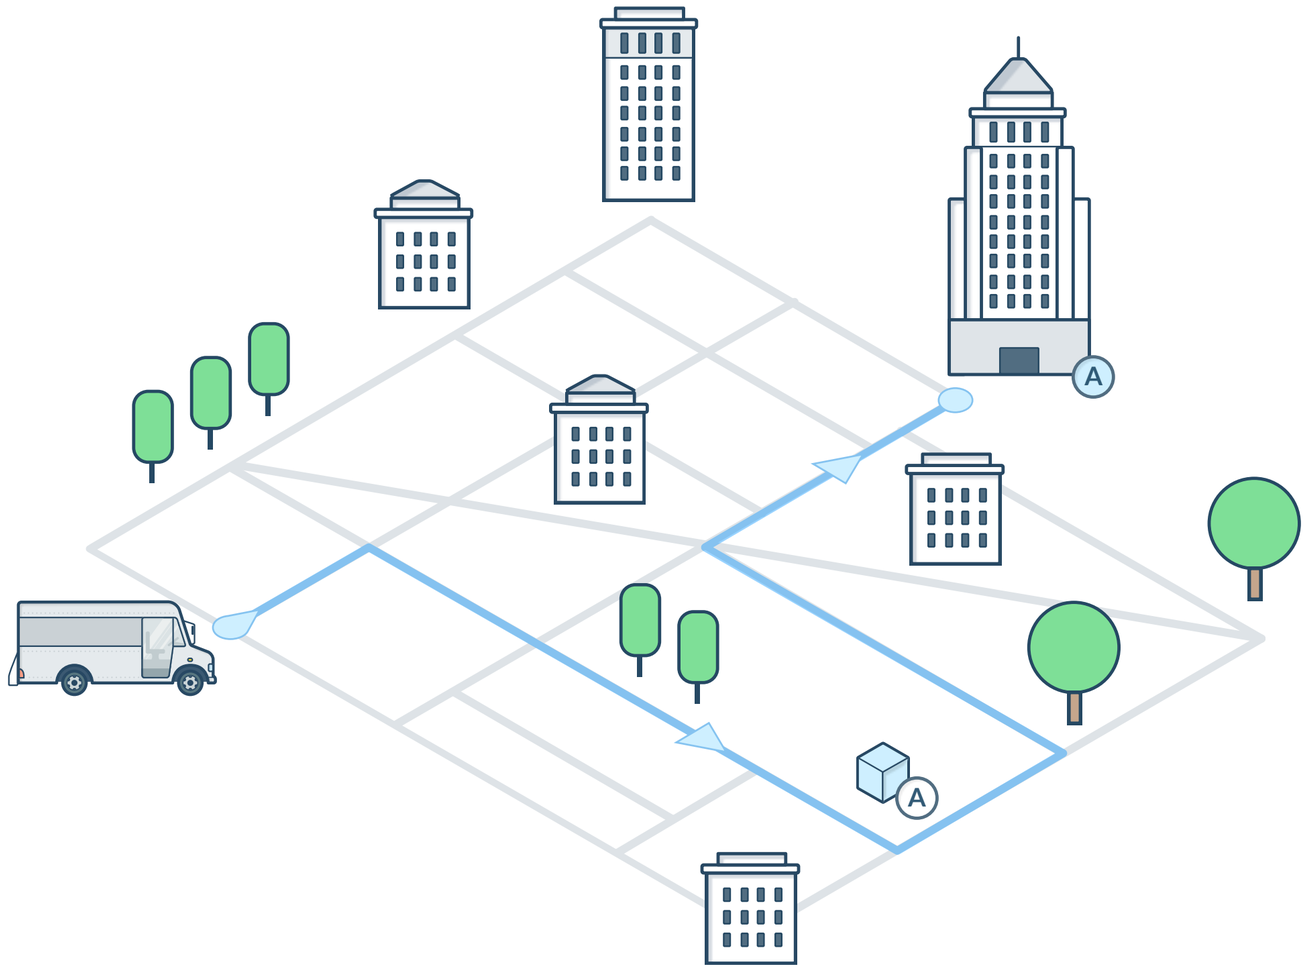

Analysis:

While this is the most straight-forward and widely-adopted approach, it is not the most effective. The method is far too simple, and because drivers are only assigned single packages, the vehicles remain heavily underutilized. For example, if you have two packages that need to be picked up in the same area, this method would require two vehicles to be dispatched.

Average score: **$10,733**

2. Greedy car pool

Method:

This approach is very similar to the previous approach in that the closest available vehicle is dispatched. The only difference is that each vehicle will pick up multiple packages at the pickup location, and deliver them one-by-one, in order of proximity.

Analysis:

This is still a very simple algorithm to implement, but performs far better than the 'Pure Greedy' approach. Efficiencies are gained because many packages are being picked up at the same location by one driver. A good example of this is when the pickup location is an office building where many packages need to be transported at once. It should be noted that this approach does not always fare well in the real world, since package pickups are not always so conveniently consolidated.

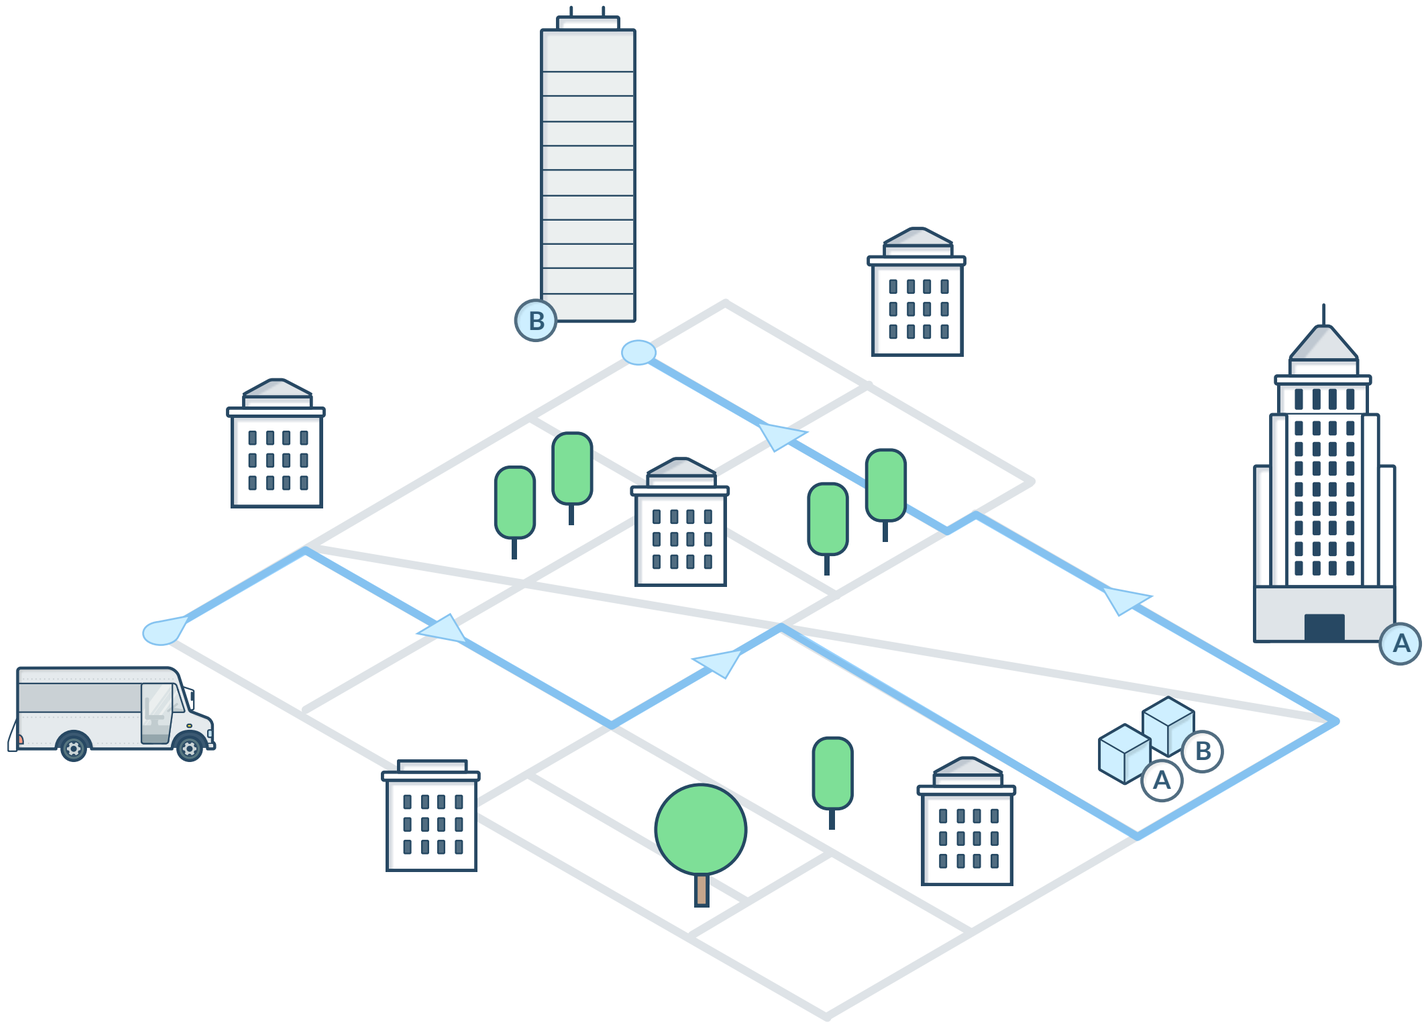

Average score: **$13,300**

3. Along the way

Method:

This approach considers not only idle vehicles, but also vehicles that are already en route to drop off a package. If a new pickup order comes in that happens to be along the route of one of the vehicles, then it will be assigned to that driver.

This may resemble a rudimentary approach to something like UberPool or Lyft Line, where drivers can pick up more passengers that are heading in roughly the same direction.

Analysis:

While this method considers the entire fleet, it only considers the rough direction each vehicle is heading towards. Drivers are still dispatched one-by-one. The disadvantage of this is that it could lead to one driver committing to a cross-town route. In the meantime, new short trips can be requested right around this driver's location, but since it isn't "along the way" it will be ignored. As you will see later, a route optimization solution would be clever enough to avoid such situations.

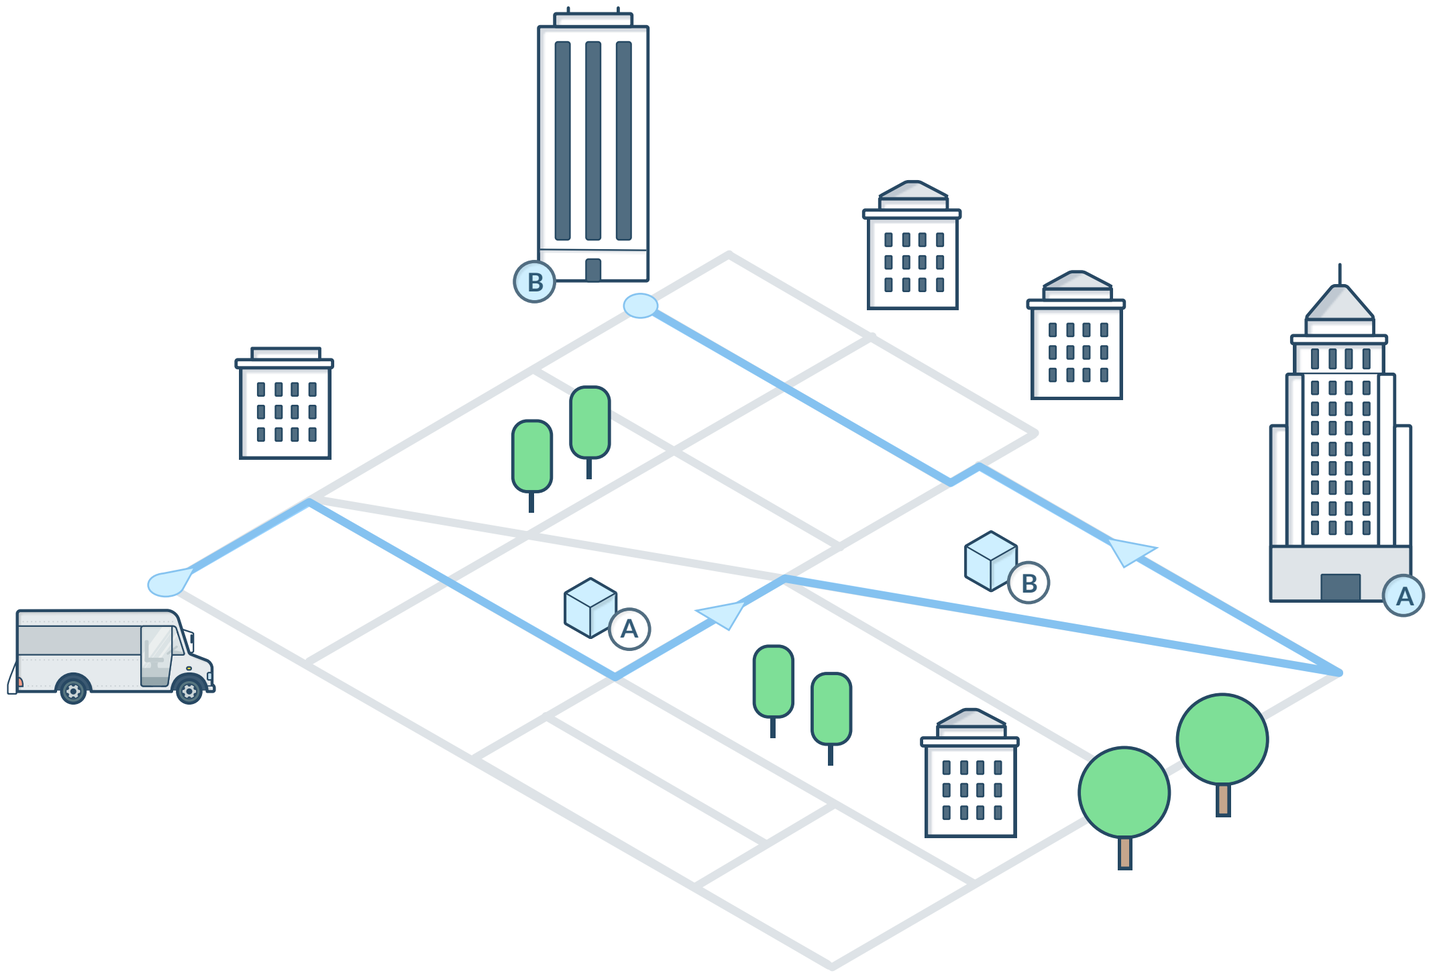

Average score: **$11,500**

4. Easy wins

Method:
This method might sound familiar, because it's what happens when old-school cab drivers deny you a ride if you request one that's too far for their liking. It is more lucrative for a cabbie to make a bunch of quick trips, because of the minimum fixed fee they earn each time.

This algorithm will look for quick and easy wins, i.e. short trips in close proximity. While driving towards a pickup location, the car can be rerouted to a pick up new package if it is an easier win. It will also pick up multiple packages at the same location.

In order to maximize the chances of finding these short trips, some cars are assigned to roam dense areas like a city's downtown core, for example.

Analysis:
This is a more sophisticated method involving demand prediction and analysis of historical trip data to discover dense and popular areas. While it may seem like a good idea to assign vehicles to certain areas, by doing so you are arbitrarily limiting the potential of your fleet.

Also, denying your customers a ride is obviously not a good idea in general. The good news is that this strategy didn't perform so well in the simulator (earning less money than the other methods) so there really is no good reason to try this at home.

Average score: **$11,667**

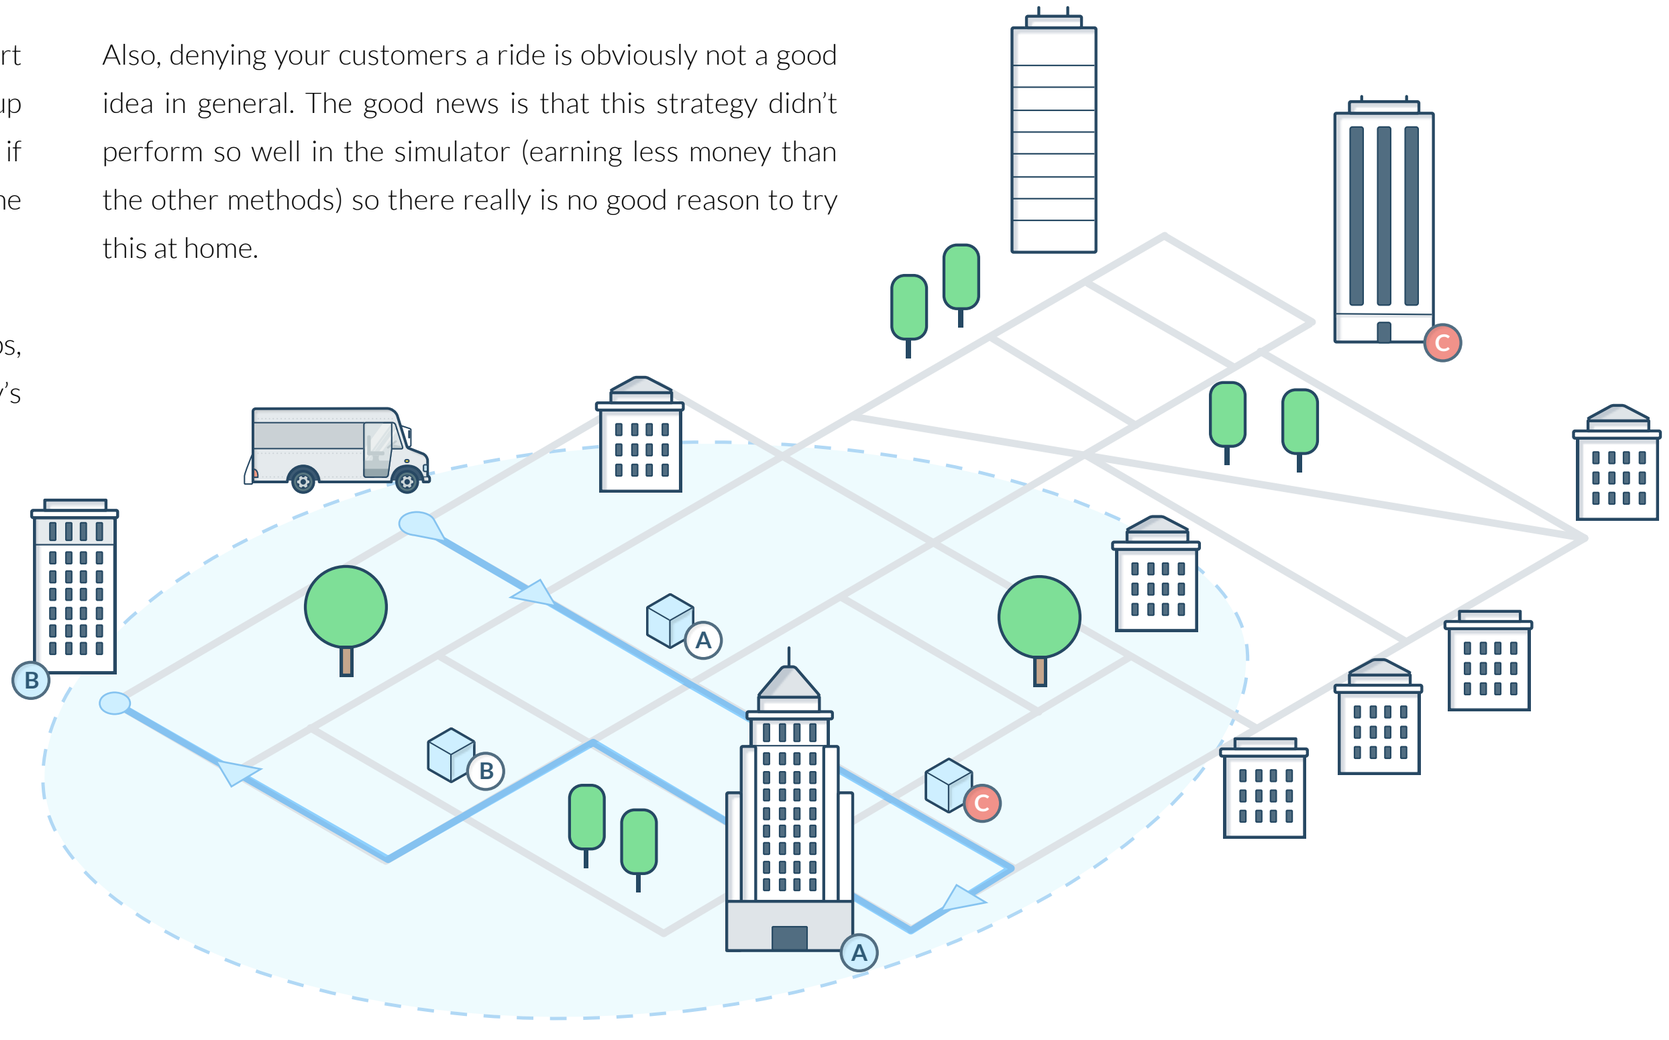

5. Defer and consolidate

Method:
The essence of this approach is to defer departure and and pick up as many packages as possible, while ensuring that none of the packages will be dropped off late. The goal is to get as much consolidation as possible and fill up the vehicle, before starting to deliver.

Analysis:
Waiting around is never a good idea. Short of predicting the future, you really don't know when and where new orders will be coming in, so you're really taking a gamble by deferring departure and hoping for some kind of consolidation. That time is much better spent dropping off your package so that you are ready again to pick up again.

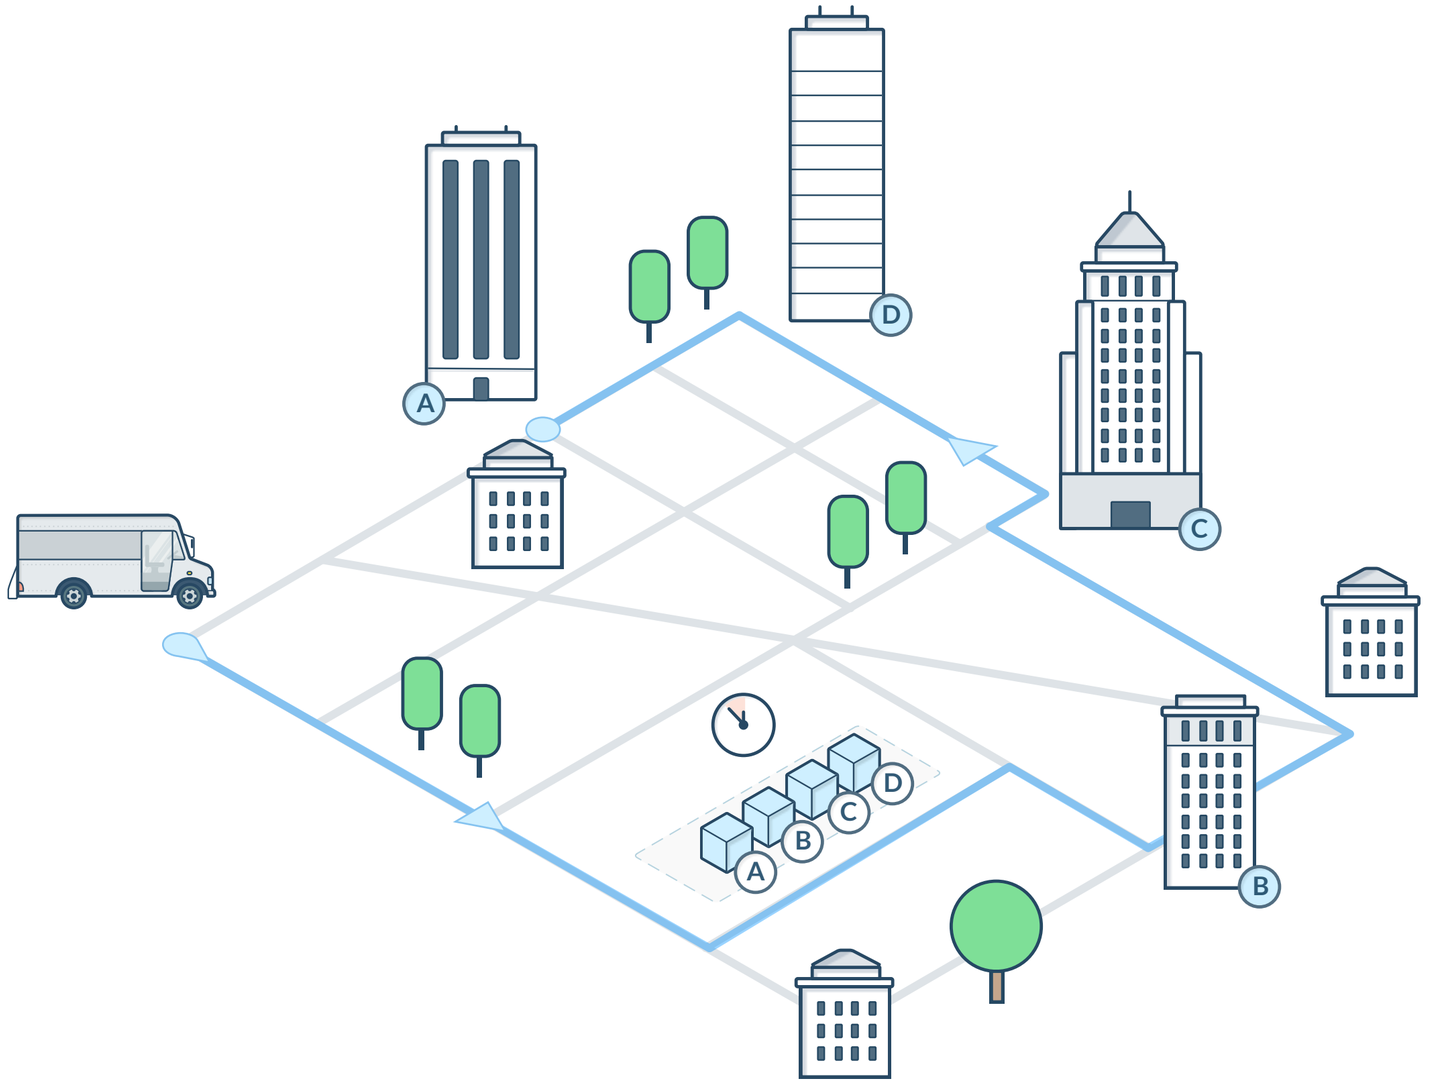

Average score: **$11,083**

6. Route optimization

Average score: $15,167

Method:

Route optimization is about finding the shortest total driving time, given a fleet of vehicles and a host of orders with numerous constraints. This is also known as the Vehicle Routing Problem. In this particular case, we need to solve a same-day pickup and delivery problem with time-windows and car capacities.

The key here is that each snapshot in time can be considered a routing problem that we solve for, the results of which become the updated driving instructions for the fleet.

To this end, we hooked up Routific's same-day pickup and delivery API.[3] As new information becomes available throughout the simulation, the entire fleet is re-optimized so they are always taking the most optimal routes.

Analysis:

By taking into account all the information that is at hand, and using sophisticated route optimization methods, enormous efficiencies can be found. Unlike previous methods, route optimization considers all the current pickup and delivery orders with their constraints, as well as the entire fleet's routes, their current and future locations, and capacities.

When new information becomes available (a new order or a cancelled order, for example) the entire fleet's routes can be re-optimized in real-time. You are continuously making the best decisions.

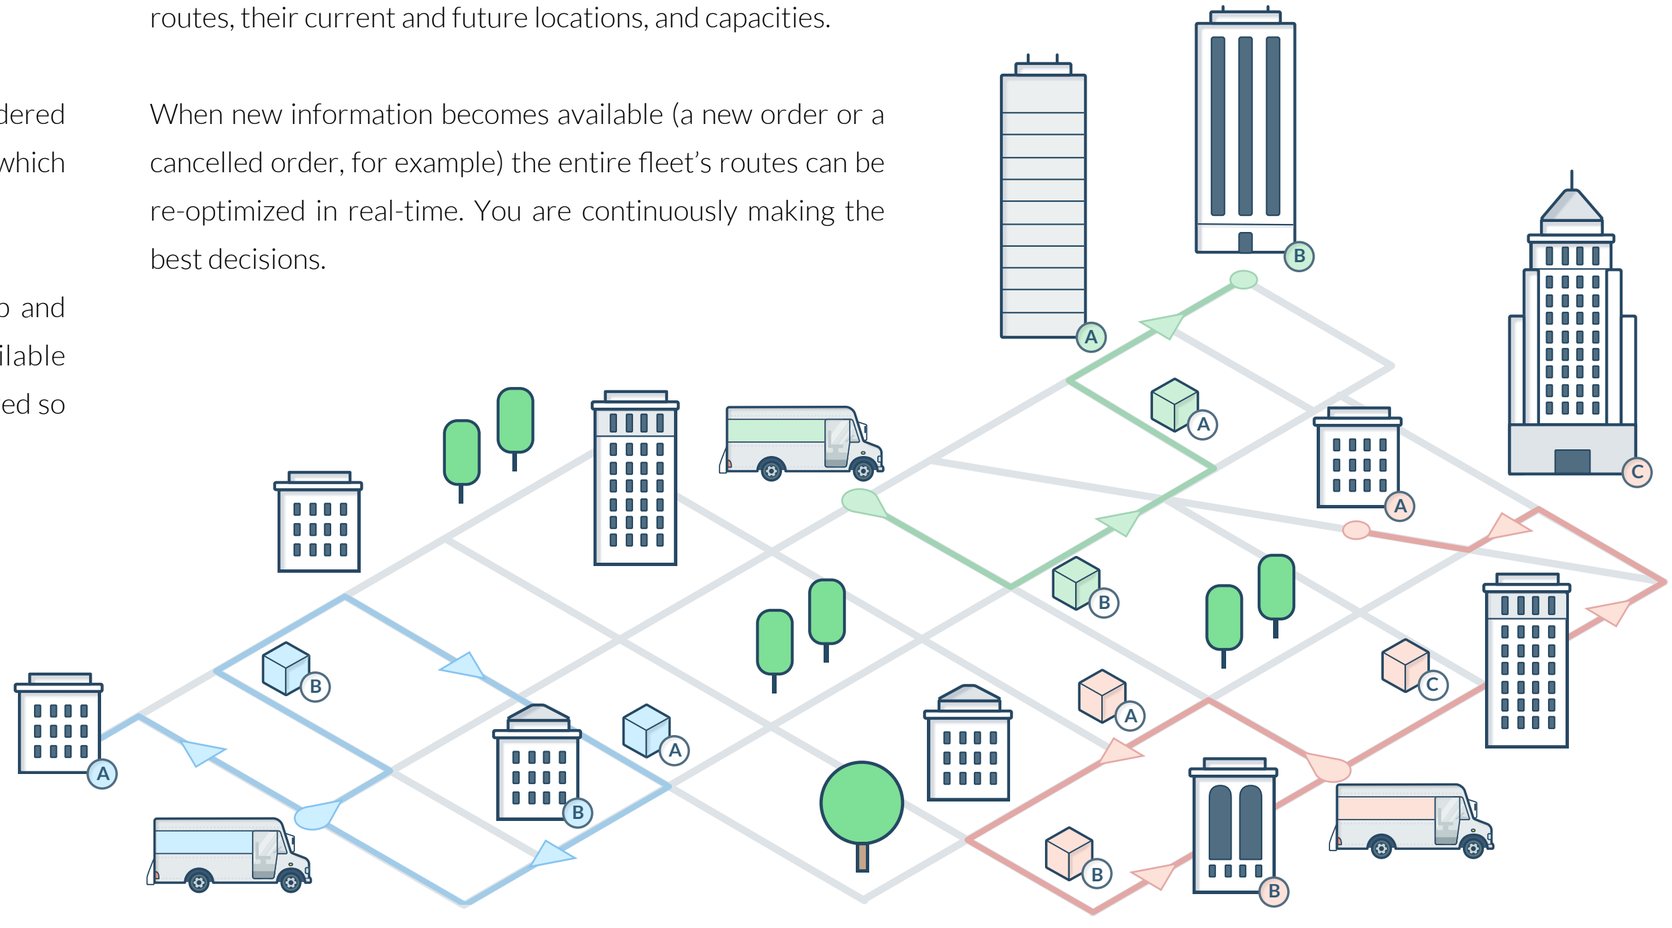

[3] See API documentation: docs.routific.com

Results

Each of the above algorithms were run three times on the simulation model. Here are the summarized results:

Route optimization clearly shines. It performs on average 30% better when compared with other approaches. When sized up against the 'Pure Greedy' approach – the simplest and most popular approach – route optimization performs 41% better.

What is surprising is that the "Greedy carpool" approach actually performs much better than the more sophisticated heuristics. But it's also important to note that in reality, packages can spawn from any location, not just at select buildings, which means that a Greedy carpool approach typically won't have as many opportunities to pick up more than one package, as it did in this simplified simulation environment.

Most companies that want to replace manual dispatchers with an auto-dispatch algorithm will implement a 'Pure Greedy' approach, because it is the easiest to build and building a route optimization solution in-house is far too complicated for most.[4] Today, businesses can leverage route optimization solutions using dedicated APIs, making it easier than ever before to replace and outperform manual dispatchers.

#	Algorithm	Run 1	Run 2	Run 3	Average
1	Pure greedy	$10,600	$10,700	$10,900	$10,733
2	Greedy carpool	$12,800	$13,650	$13,450	$13,300
3	Along the way	$11,000	$11,850	$11,650	$11,500
4	Easy wins	$12,150	$11,750	$11,100	$11,667
5	Defer and consolidate	$10,300	$11,100	$11,850	$11,083
6	Route optimization	$14,650	$15,250	$15,600	**$15,167**

[4] Frankly, building anything in-house that is otherwise widely available as an API is strongly advised against. Read about both sides of the argument in this article: www.medium.com/@kuomarc/don-t-build-it-in-house-there-is-an-api-for-that-c929b8677137

Why Route Optimization Reigns

A route optimization solution allows an entire fleet to operate as one centralized system. It considers all the information at hand, and can deliver routes that include multiple pickups and dropoffs intertwined.

The other heuristics are only capable of considering one package at a time. Route optimization is the only method that considers both the packages already on the vehicles, as well as every outstanding package yet to be picked up. The entire fleet then decides together, as a collective, who picks up what in the most efficient way possible.

The routing problem is known to become exponentially more difficult with larger problems.[5] So while route optimization already outperforms the other heuristics in this simplified simulation environment by 30%, it will only shine more in a more complex real-world scenario.

We have already known this to be the case with scheduled routing, where route optimization can find solutions that are 37% shorter.[6] We can now also conclude that route optimization brings equally significant efficiencies to on-demand scenarios.

[5] The number of possible permutations of routes is n! (factorial). So with 5 packages, you already have 120 possibilities to consider. When you have 57 packages to dispatch, you already have 40.5 quattuorvigintillion possible routes to consider! The only way to tackle that is with route optimization.

[6] For this and more case studies go to: www.routific.com/stories/

DRYV Case study

How one company delivers in an on-demand world

DRYV is an on-demand laundry and dry cleaning service, currently operating in 3 markets: Chicago, Detroit and Los Angeles.

DRYV partners with high quality retail cleaners, offering pickup and delivery services on their behalf. They also work with property management companies and businesses to ensure people have access to a convenient and affordable laundry and dry cleaning solution whether they're at home or at work.

DRYV Case study

Business Challenges:

Building and scaling a delivery fleet is capital-intensive and extremely complex. The team behind DRYV collected loads of data to understand the demands of the market and worked on building cutting-edge technical and logistical infrastructure to support it.

Route planning and dispatching was purely a manual task, at least in the beginning.

“We had someone sit in front of a computer, plot the stops on the map and drag things around to create a route,” said Dan Parsons, co-founder and co-CEO of DRYV. “That’s fine when you have one driver, but it soon became way too complex and we knew we were spitballing it.”

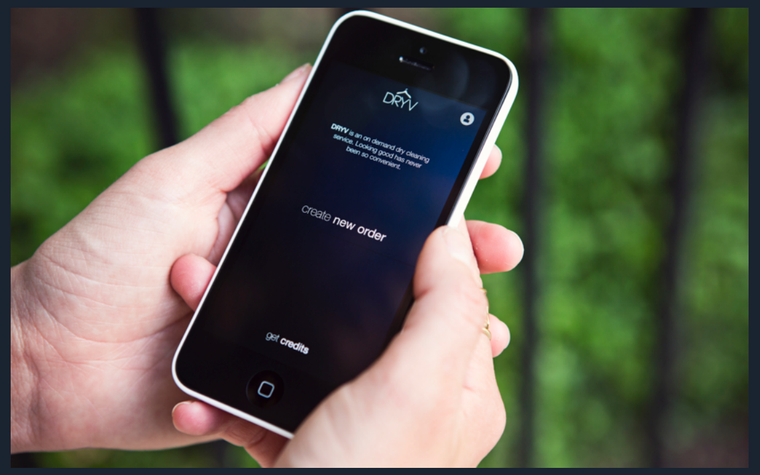

Solution:

DRYV integrated Routific’s API to automate and optimize their dispatching operations.

“The API was very straightforward and easy to integrate. Everything was well documented and the team behind Routific provided us with superb support,” Parsons said. “For such a complex piece of technology, integration was a quick and painless experience.”

Routific also completely eliminated the need for manual dispatchers. DRYV simply calls the API throughout the day and receives optimized routes in return. When last-minute orders and changes to orders come in, DRYV can rely on the API to respond quickly with a new result.

> “ Our mission is to leverage technology and to pass that value to both the retail cleaners we partner with and the consumers that trust us to get their laundry delivered clean and on time. Routific is essential in this process.”
>
> — Dan Parsons, co-founder and co-CEO of DRYV