Global Services

LP in industry – Airline planning and operations

Mattias Grönkvist
Optimization Chapter Lead, Jeppesen Crew & Ops

February, 2019

Agenda

Global Services

- Introduction to Jeppesen
- Airline planning and operations
- Schedule and fleet planning problems
- Crew planning problems
- Operational problems
- Follow-up

The Boeing Company

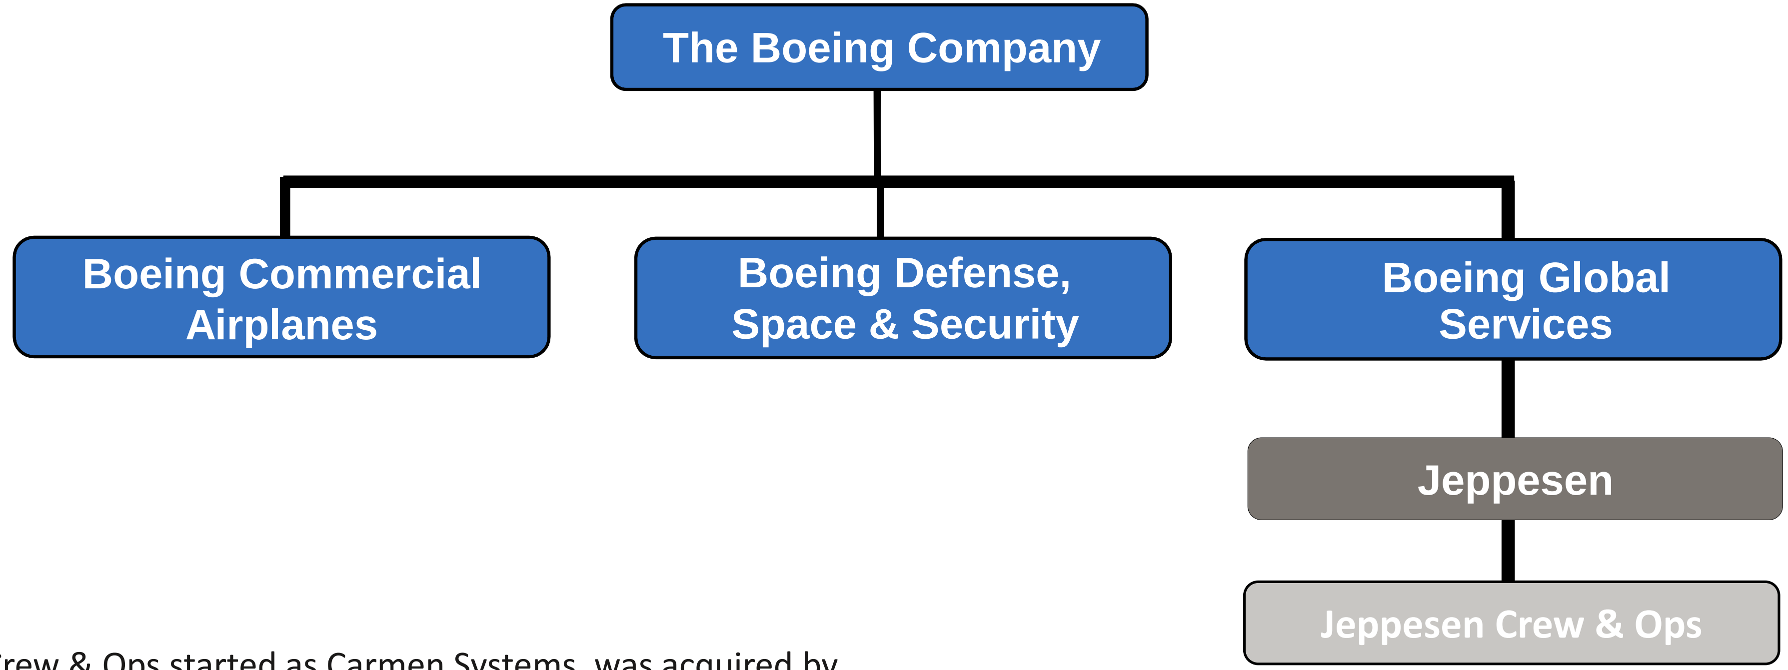

Jeppesen Crew & Ops started as Carmen Systems, was acquired by Jeppesen in 2006.

We are about 350 employees in the Gothenburg office.

What does Jeppesen Crew & Ops do?

We create market-leading software to model, solve and analyze planning and operations problems for airlines

Airline planning problems are some of the largest, hardest and most well-studied real-world optimization problems

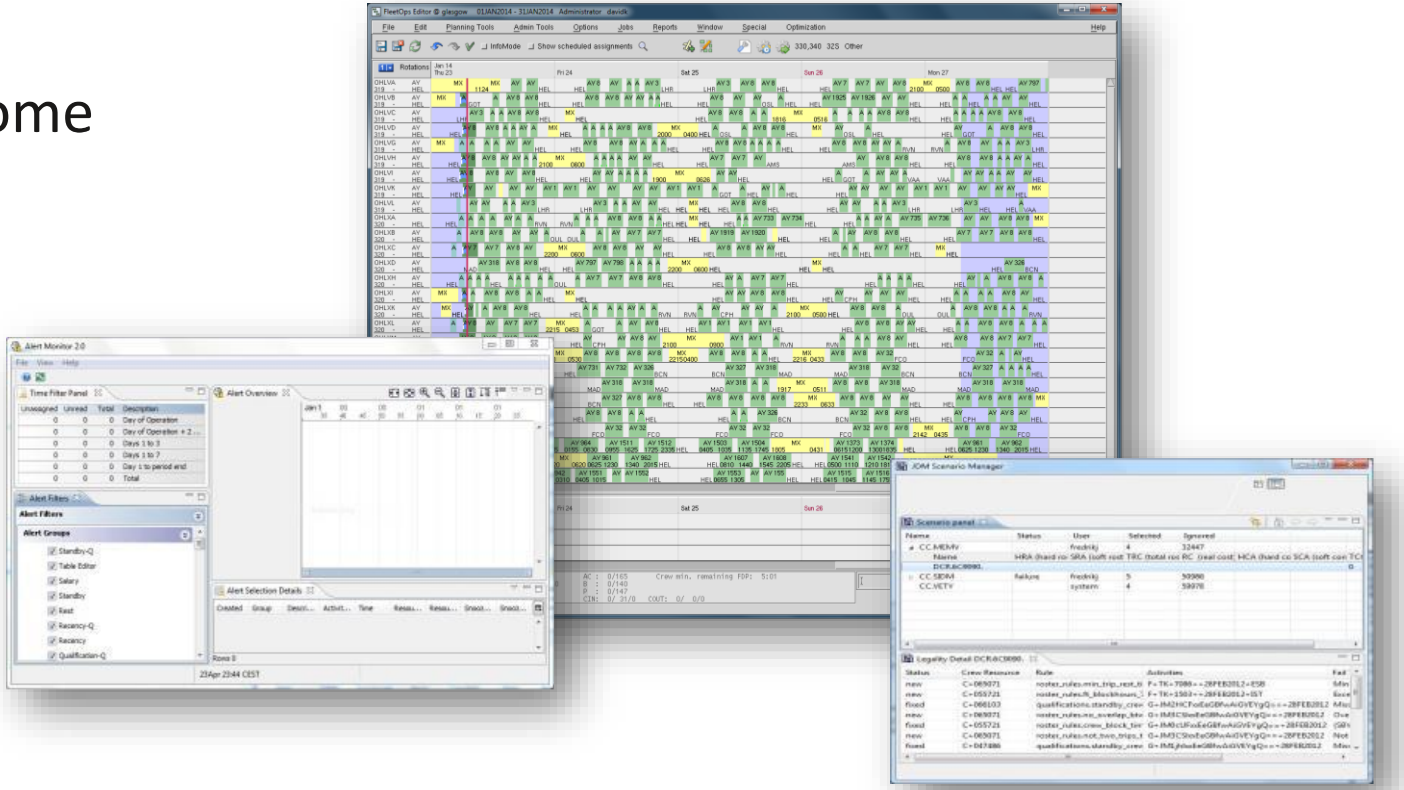

Jeppesen Crew & Ops Customers

Global Services

Aer Lingus
AEROMEXICO
AEROMEXICO connect
AIR CANADA
AIR NEW ZEALAND
AIRFRANCE
Alaska Airlines
Alitalia
Avianca
BRITISH AIRWAYS
DB BAHN

DELTA
DSB
Emirates
EXPRESSJET
EVA AIR
FINNAIR
green cargo
GOL Linhas aéreas inteligentes
IBERIA
interjet
jetBlue

KLM
KLM Cityhopper
LAN TAM
LATAM AIRLINES GROUP
Lufthansa
Philippine Airlines
PEGASUS AIRLINES
QANTAS
QATAR AIRWAYS القطرية
السعودية SAUDIA
SAS Scandinavian Airlines
SINGAPORE AIRLINES

SkyWest AIRLINES
SWISS
Thomson Airways
transavia.com
TURKISH AIRLINES
UNITED
U·S AIRWAYS
widerøe
Virgin atlantic
Virgin australia
volaris

Airline planning and operations

Customer expectations on airlines

- **Point-to-point transportation**
  - I want to be transported from my origin to my destination
- **On-time performance**
  - I want to depart and arrive on the agreed times
- **Certain level of comfort**
  - I want to have the leg room and noise level I have been promised
- **Certain level of service**
  - I want to get the service I have been promised, e.g. meals and drinks
- **Maximum safety**
  - I want my transport to be as safe as possible
- **Fixed price**
  - I want to pay the agreed price

Airline resources

Global Services

- **Timetable of flights**
  - Where do we fly?
  - When do we fly?

- **Aircraft**
  - Aircraft are grouped into families/fleets
  - Aircraft have different sizes, crew needs and equipment

- **Crew**
  - Cabin crew and cockpit crew
  - Maybe several crew bases
  - Maybe several different union agreements

- **Ground staff**
  - At airports, control centers, maintenance facilities

- **Ground resources**
  - Maintenance facilities
  - Simulators
  - Course rooms

The crew and fleet planning processes

Global Services

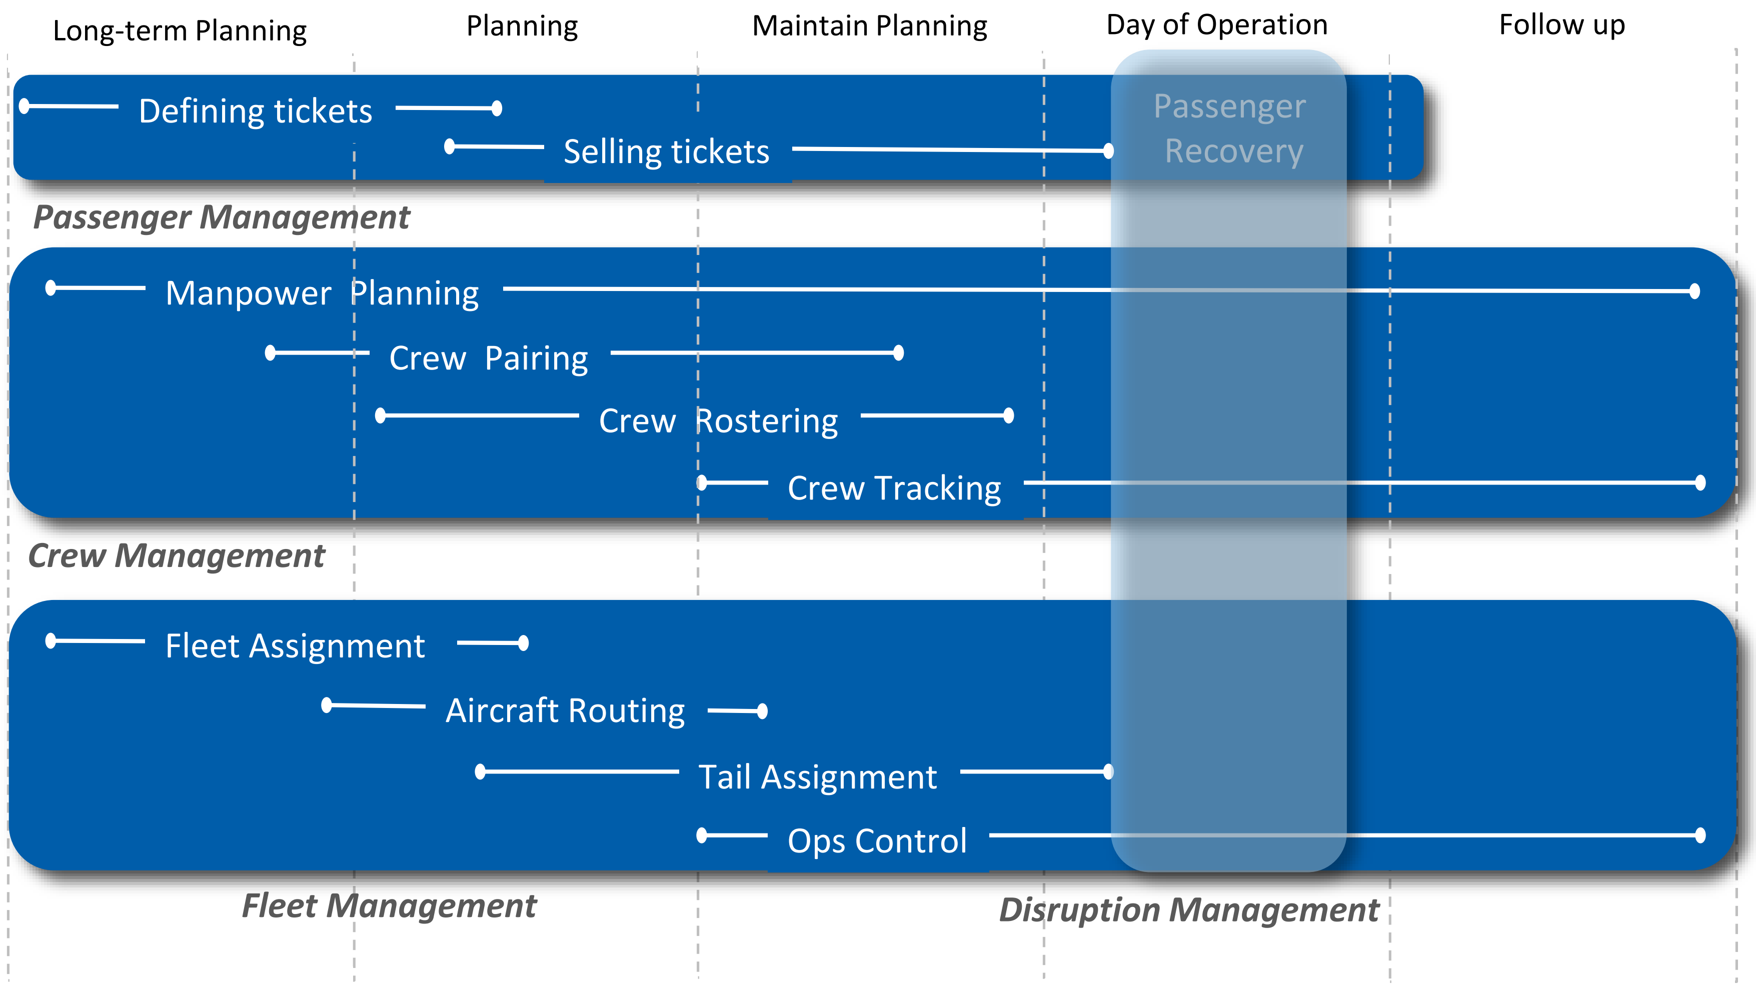

Schedule and fleet planning problems

Schedule and fleet planning

Global Services

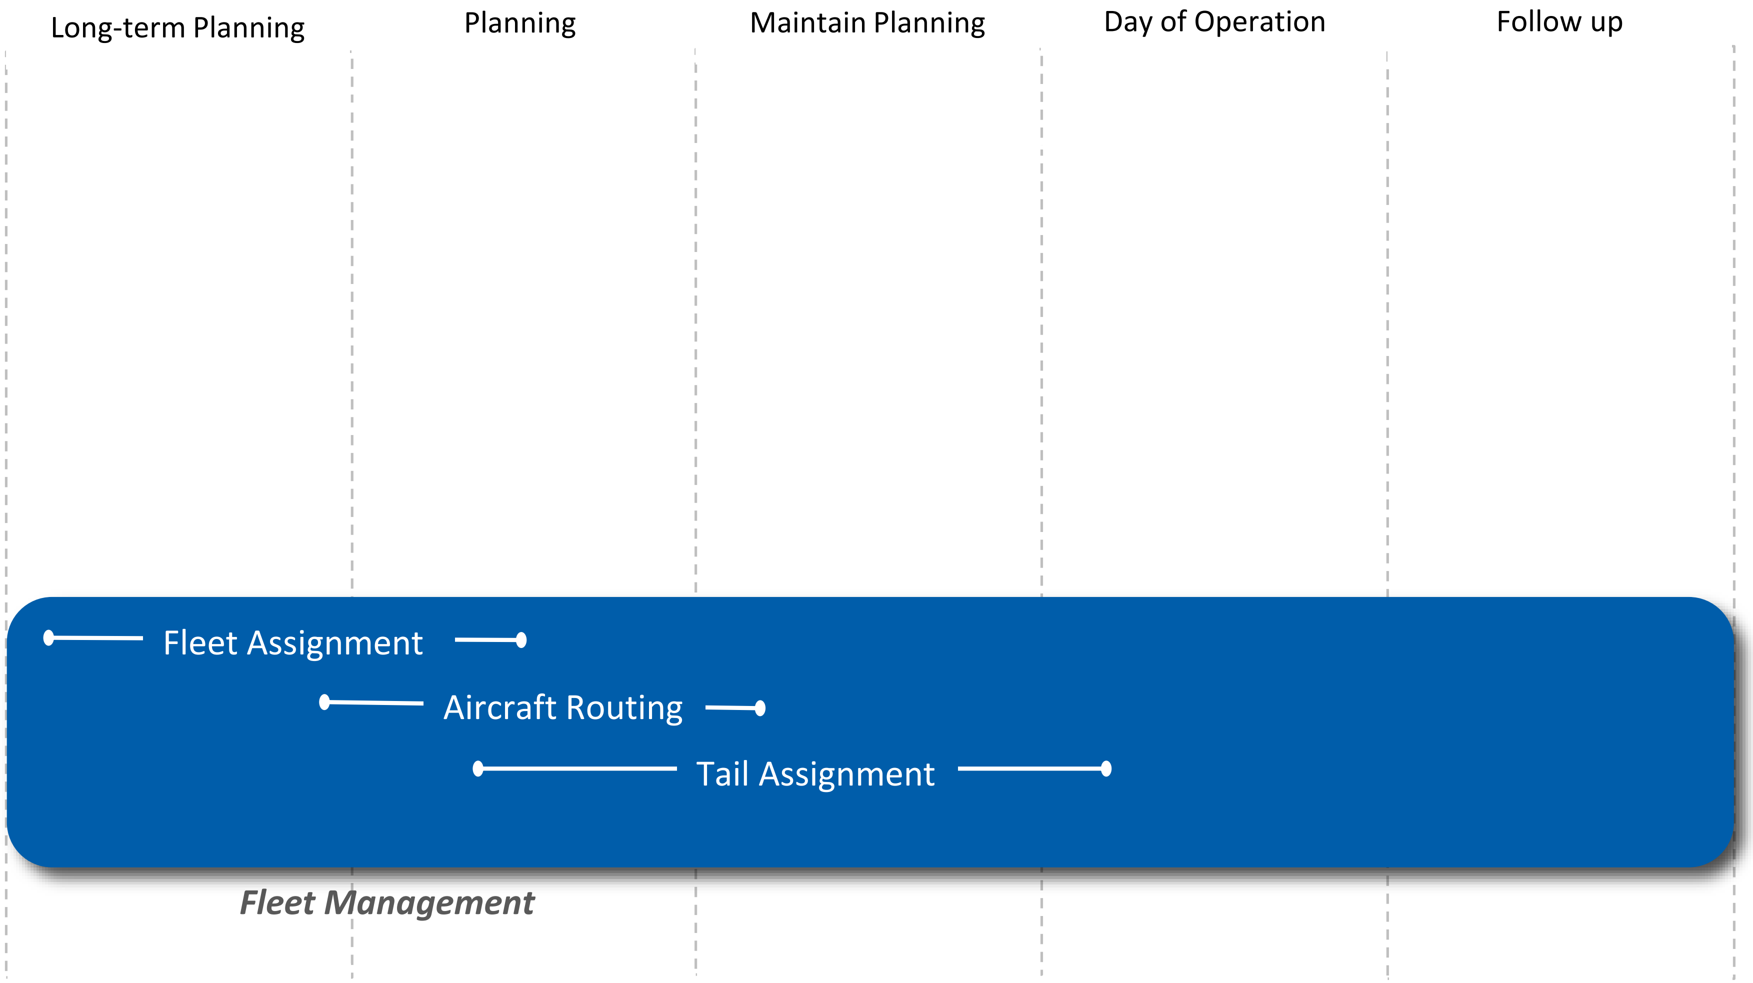

Flight scheduling

Global Services

Flight scheduling is the process of creating a timetable. A bunch of questions must be answered:

- Where do people want to fly?
- What are they willing to pay?
- How many passengers are willing to fly from X to Y for a ticket price of Z?
- Where and when do our competitors fly?

Creating a flight schedule proposal is often a manual task, but **evaluating** it is not…

The fleet assignment problem

Assign aircraft types to all flights

In such a way that

- Expected profit is maximized
- No more than the available number of aircraft is used
- Aircraft flow balance is preserved
- Fleet-level connection rules are respected

Fleet assignment is one of the most central problems for airlines.
Why?

As a consequence, it is also the most studied problem in airline OR.

Modeling the fleet assignment problem

What is the problem data?

- A set of flights *F* serving a set of airports *A*
- A set of aircraft types *T*
- Sizes of aircraft fleets for each type t, S_t
- Expected profit (cost minus revenue) for assigning aircraft type t to flight f, p_{ft}

What are suitable decision variables?

- Let's try binary variables 'flight *f* should be assigned to type *t*', x_{ft}

What is the objective?

- Maximize the total expected profit

What are the constraints?

- Exactly one aircraft type assigned to each flight
- All aircraft need to arrive at an airport before they can depart
- The total number of aircraft used of each type is at most the fleet size

Modeling the fleet assignment problem

The traditional way to model an airline networks is a *timeline network*.

Flight decision variables *x*

s show incoming and outgoing flights at an airport as

ig arcs:

Ground flow variables *g*

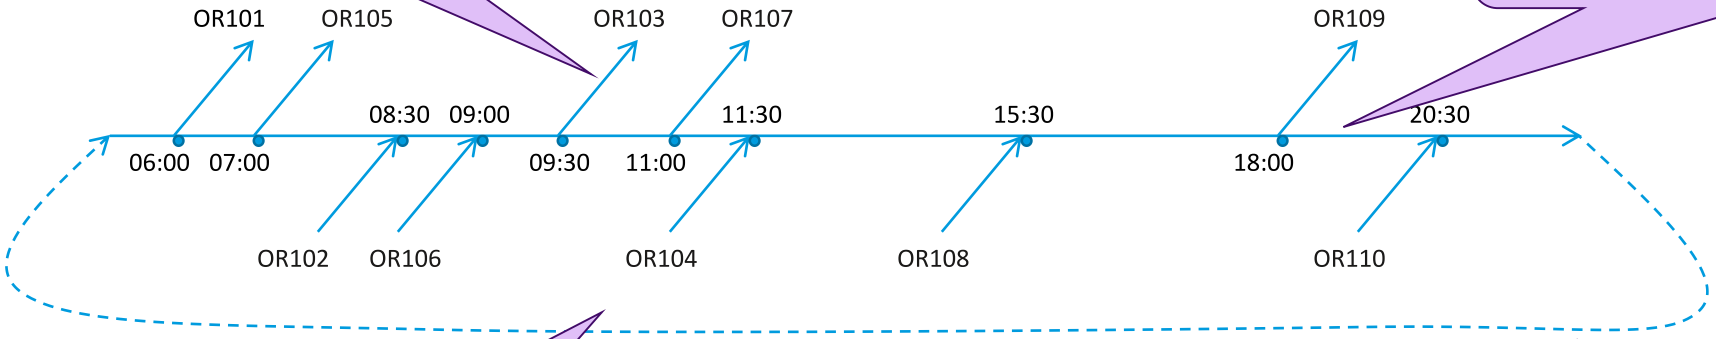

Arrival and de re *events* are represented as nodes. Backward arc represe

One network per aircraft type!

Overnight variable for airport *a*

Modeling the fleet assignment problem

Global Services

$$
\begin{aligned}
\max \quad & \sum_{f \in F} \sum_{t \in T} p_{ft} x_{ft}, \\
\text{s.t.} \quad & \sum_{t \in T} x_{ft} = 1, \quad \forall f \in F, \\
& g_{ft}^{dep\_in} = g_{ft}^{dep\_out} + x_{ft}, \quad \forall f \in F, \forall t \in T, \\
& g_{ft}^{arr\_in} + x_{ft} = g_{ft}^{arr\_out}, \quad \forall f \in F, \forall t \in T, \\
& \sum_{a \in A} g_{ta}^{overnight} \leq S_t, \quad \forall t \in T \\
& x \in \{0, 1\} \\
& g \in I
\end{aligned}
$$

Maximize the total profit

Assign a type to each flight

Flow balance at each flight arrival/departure

Only use the available number of aircraft

Decision variables are binary, ground flow variables integer

Solving the fleet assignment problem

The presented model is a **multi-commodity network flow model** – one network flow model per aircraft type, connected by a few constraints.

In the general case this problem is **NP-hard**.

Fleet assignment is typically solved using standard branch-and-bound techniques, possibly mixed with integer heuristics, local search etc.

In addition to the decision about aircraft types the model also produces routes for aircraft as a useful 'side effect'.

Fleet assignment extensions

The presented fleet assignment model is very basic. A lot of extensions and ’complications’ can be introduced:

- Change the timeline network to a more expressive but larger *connection network*.
- Introduce passenger *spill modeling* – not all passengers who want to fly, can fly.
- In addition to spill, model passenger *recapture* – some spill can be captured on other flights.
- Instead of a ’leg-based’ profit model, look at origin-destination itineraries.
- Allow flights to be re-timed slightly to gain extra profit.

The tail assignment problem

Global Services

Assign aircraft to all flights

In such a way that

- Pre-planned maintenance activities are respected
- Destination restrictions and curfews are respected
- Connection rules, possibly aircraft specific ones, are respected
- Aircraft routes allow for all necessary maintenance to be performed, e.g. every 6 days, 100 flying hours, …
- Some objective is optimized (fuel costs, robustness)

The tail assignment problem

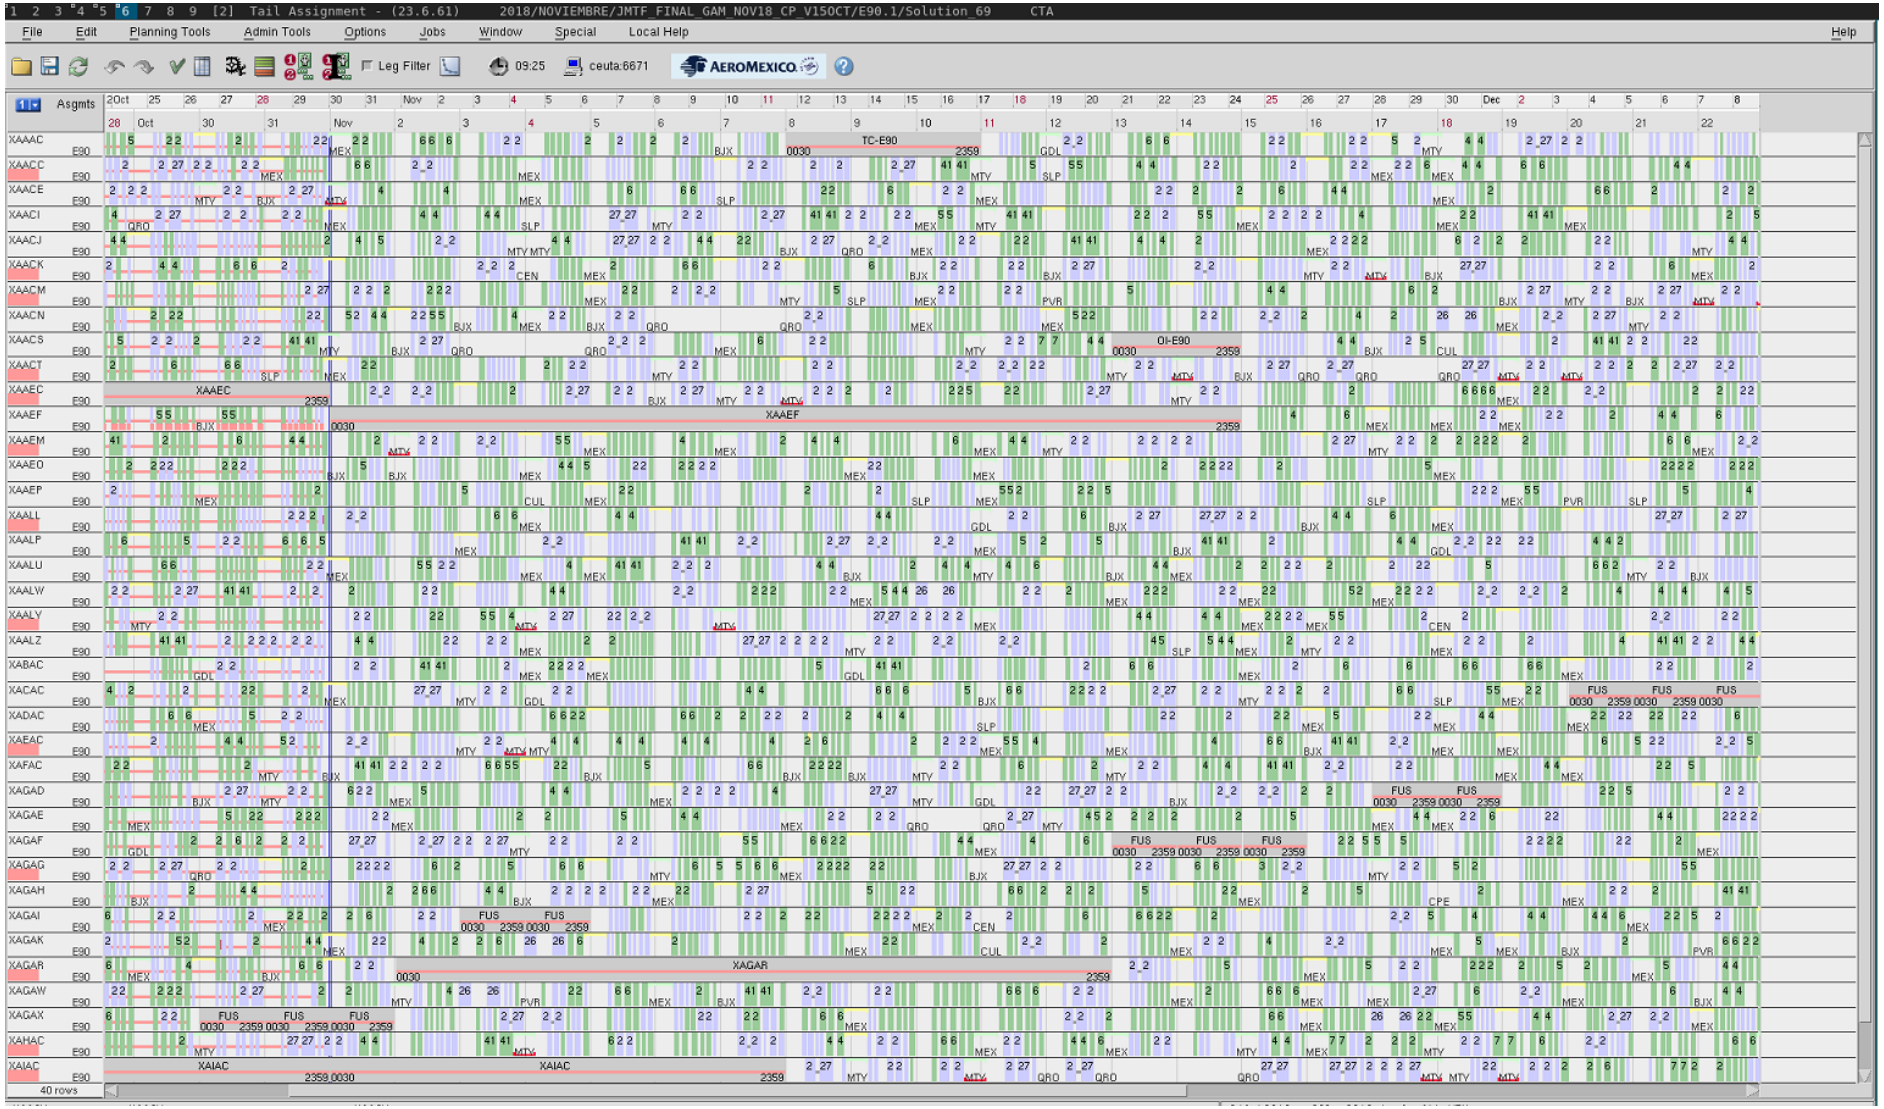

Modeling the tail assignment problem

For tail assignment it is not straightforward to use decision variables like 'flight *f* is assigned to aircraft *a*'. **Why not?**

To simplify modeling of maintenance rules, it makes sense to use **aircraft routes as decision variables**. This is an example of so-called *Dantzig-Wolfe Decomposition*.

Dantzig-Wolfe decomposition separates concerns into **generation of routes** and **selection of routes**.

What is the main problem with this choice of decision variables?

The number of variables goes from polynomial to exponential!

Modeling the tail assignment problem

Global Services

c_r is the cost of using route r

R is the set of all possible individual aircraft routes

$$\min \sum_{r \in R} c_r x_r,$$

$$\text{s.t.} \sum_{r \in R} a_{fr} x_r = 1, \quad \forall f \in F,$$

$$\sum_{r \in R} b_{gr} x_r \leq d_g, \quad \forall g \in G,$$

$$x_r \in \{0, 1\}$$

F is the set of flights

a_{fr} is *1* if route *r* covers flight *f*, *0* otherwise

G is the set of custom global constraints

b_{gr} is the contribution to global constraint *g* from route *r*

x_r are binary decision variables over *R*

Modeling the tail assignment problem

Each column represents a route

Each row represents a flight

$$\begin{bmatrix} 1 & 1 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 1 & 0 & \dots & 0 \\ 1 & 1 & 1 & 0 & 1 & \dots & 1 \\ 0 & 1 & 0 & 1 & 0 & \dots & 1 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 1 & 0 & 0 & 1 & 1 & \dots & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ \dots \\ x_n \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ \dots \\ 1 \end{bmatrix}$$

This matrix is far too large to generate – how can we restrict ourselves to interesting routes?

Solving the tail assignment problem – column generation

Global Services

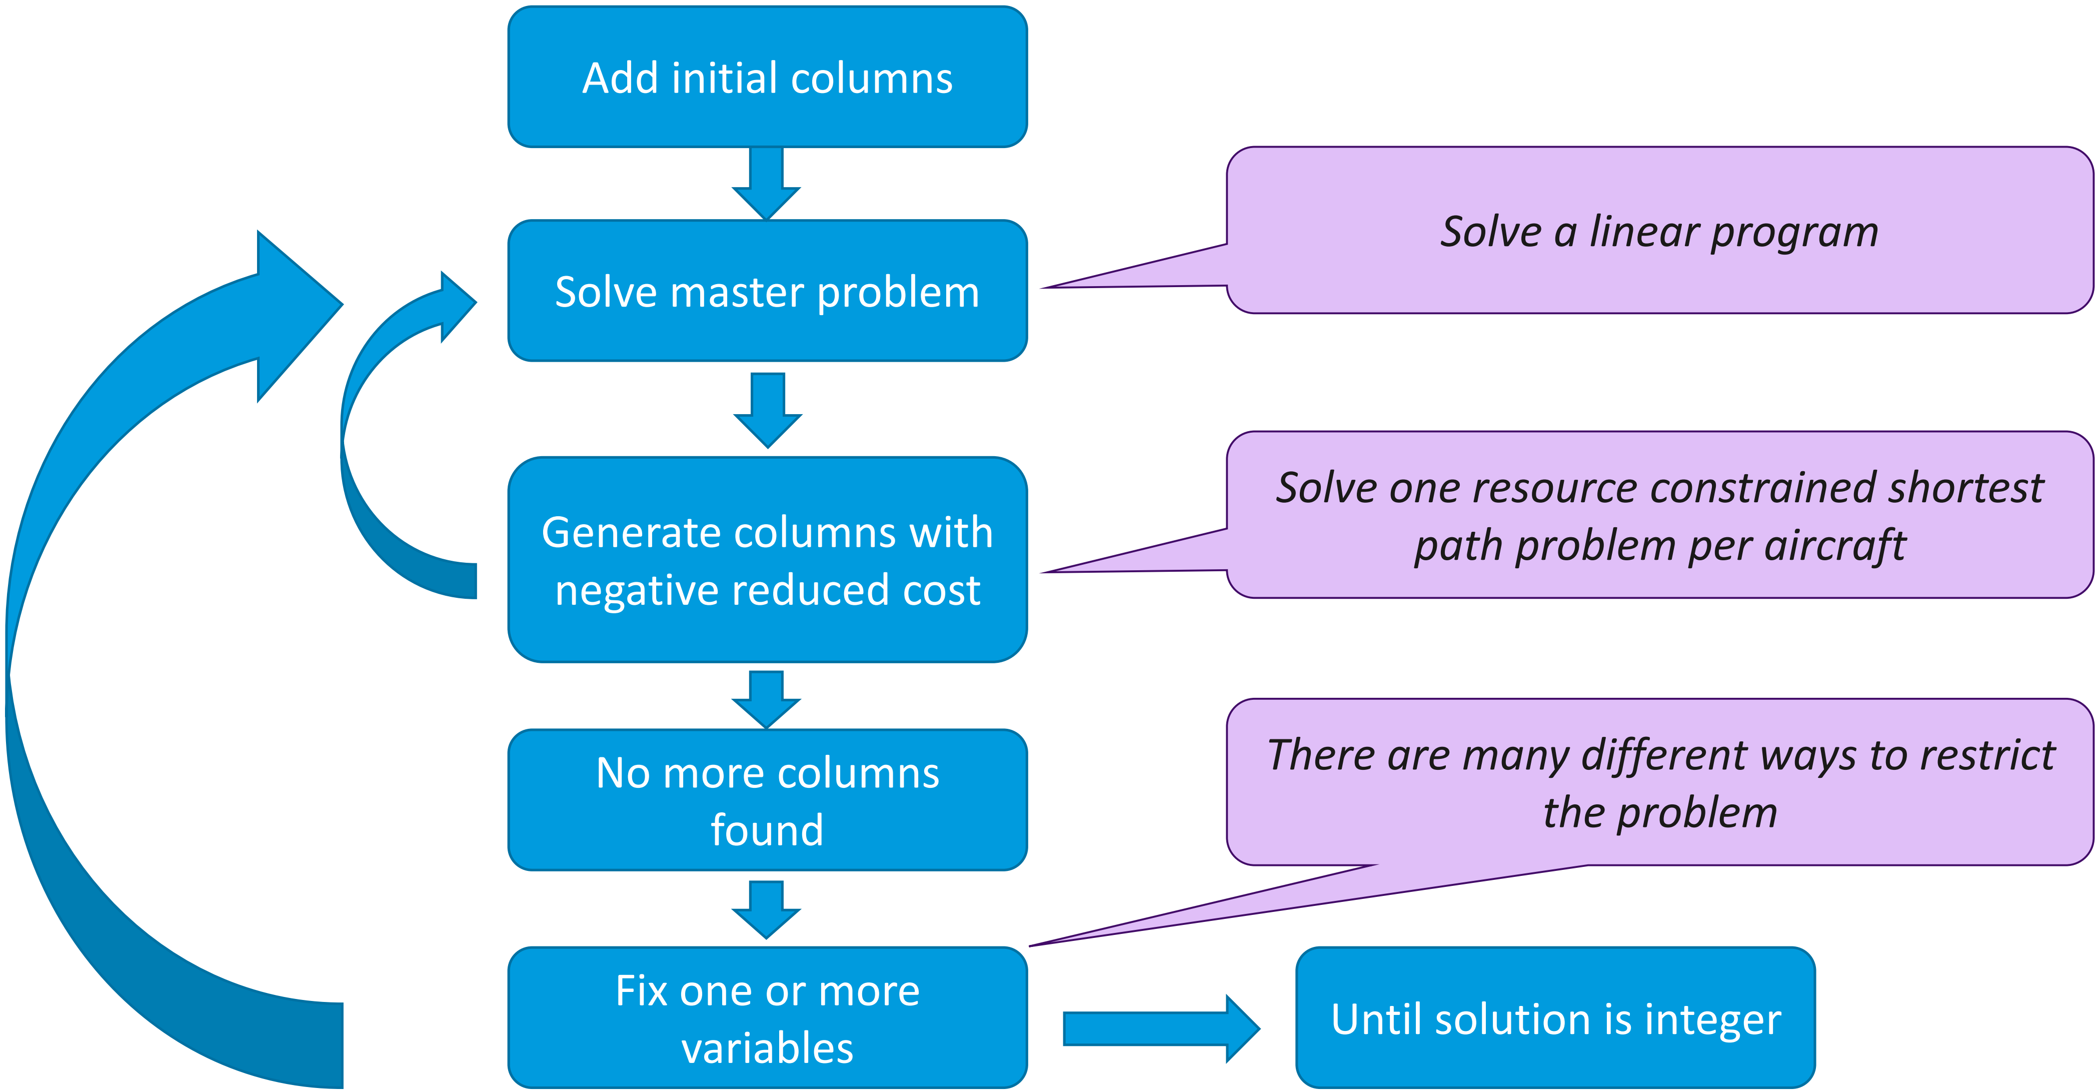

Crew planning problems

The crew and fleet planning processes

Global Services

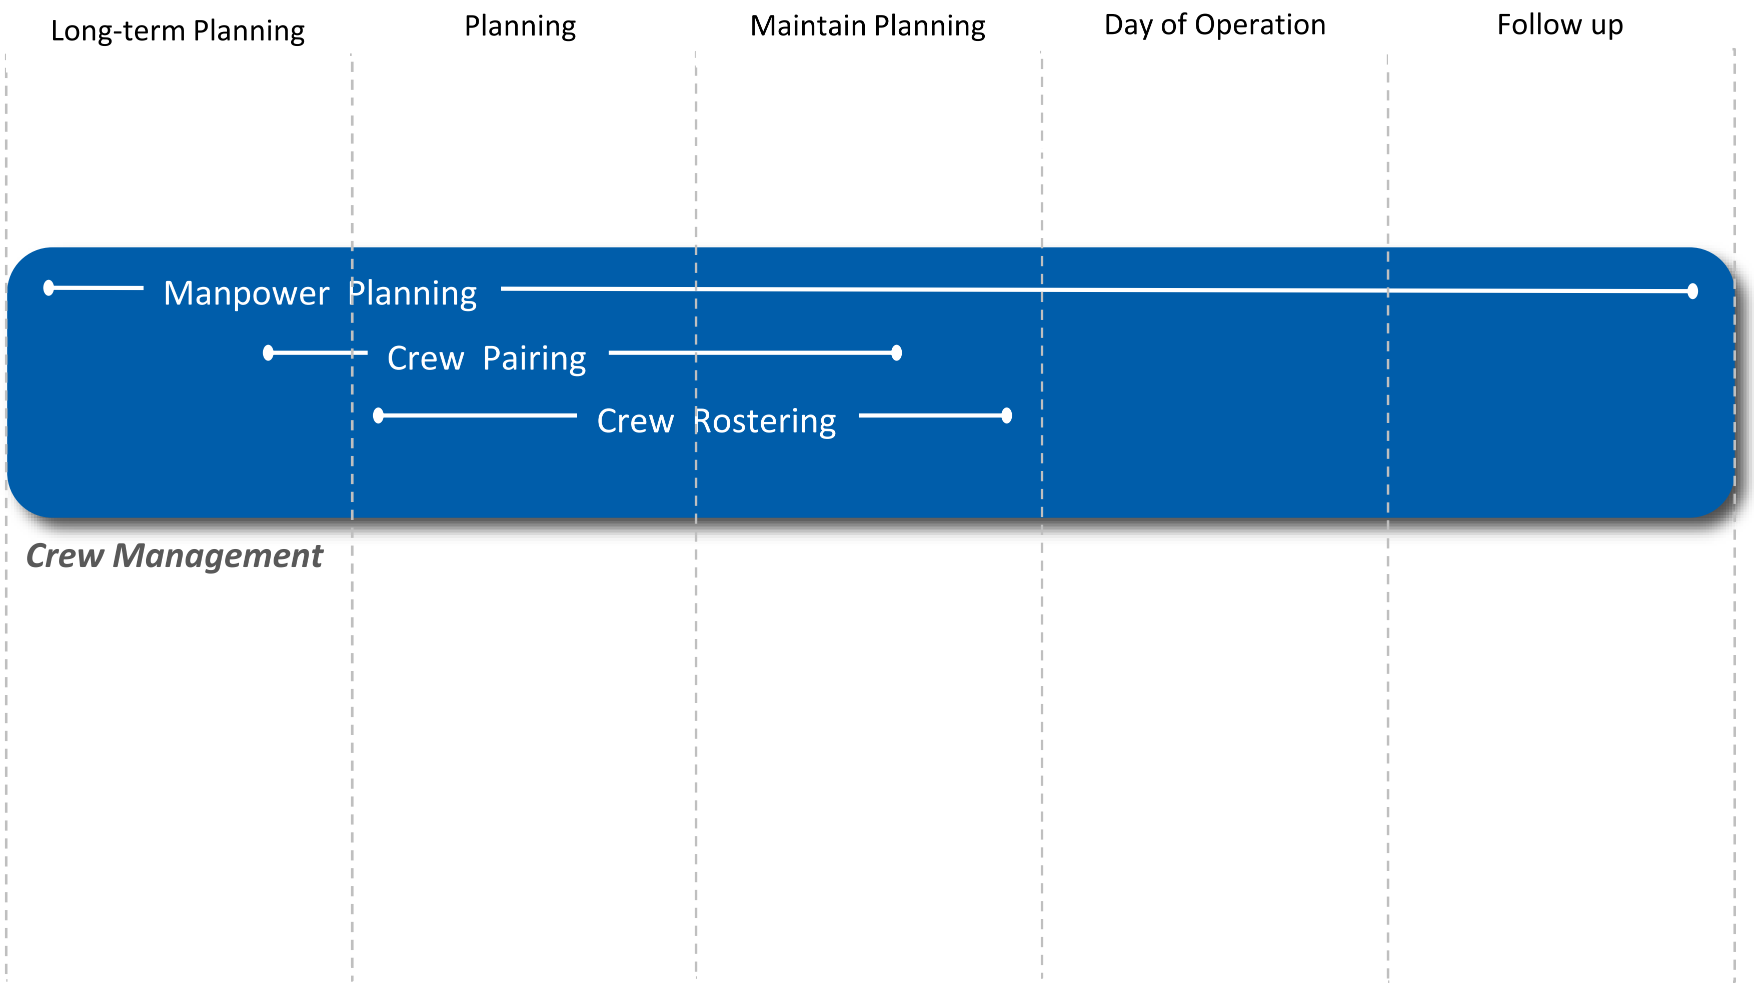

The manpower planning problem

Global Services

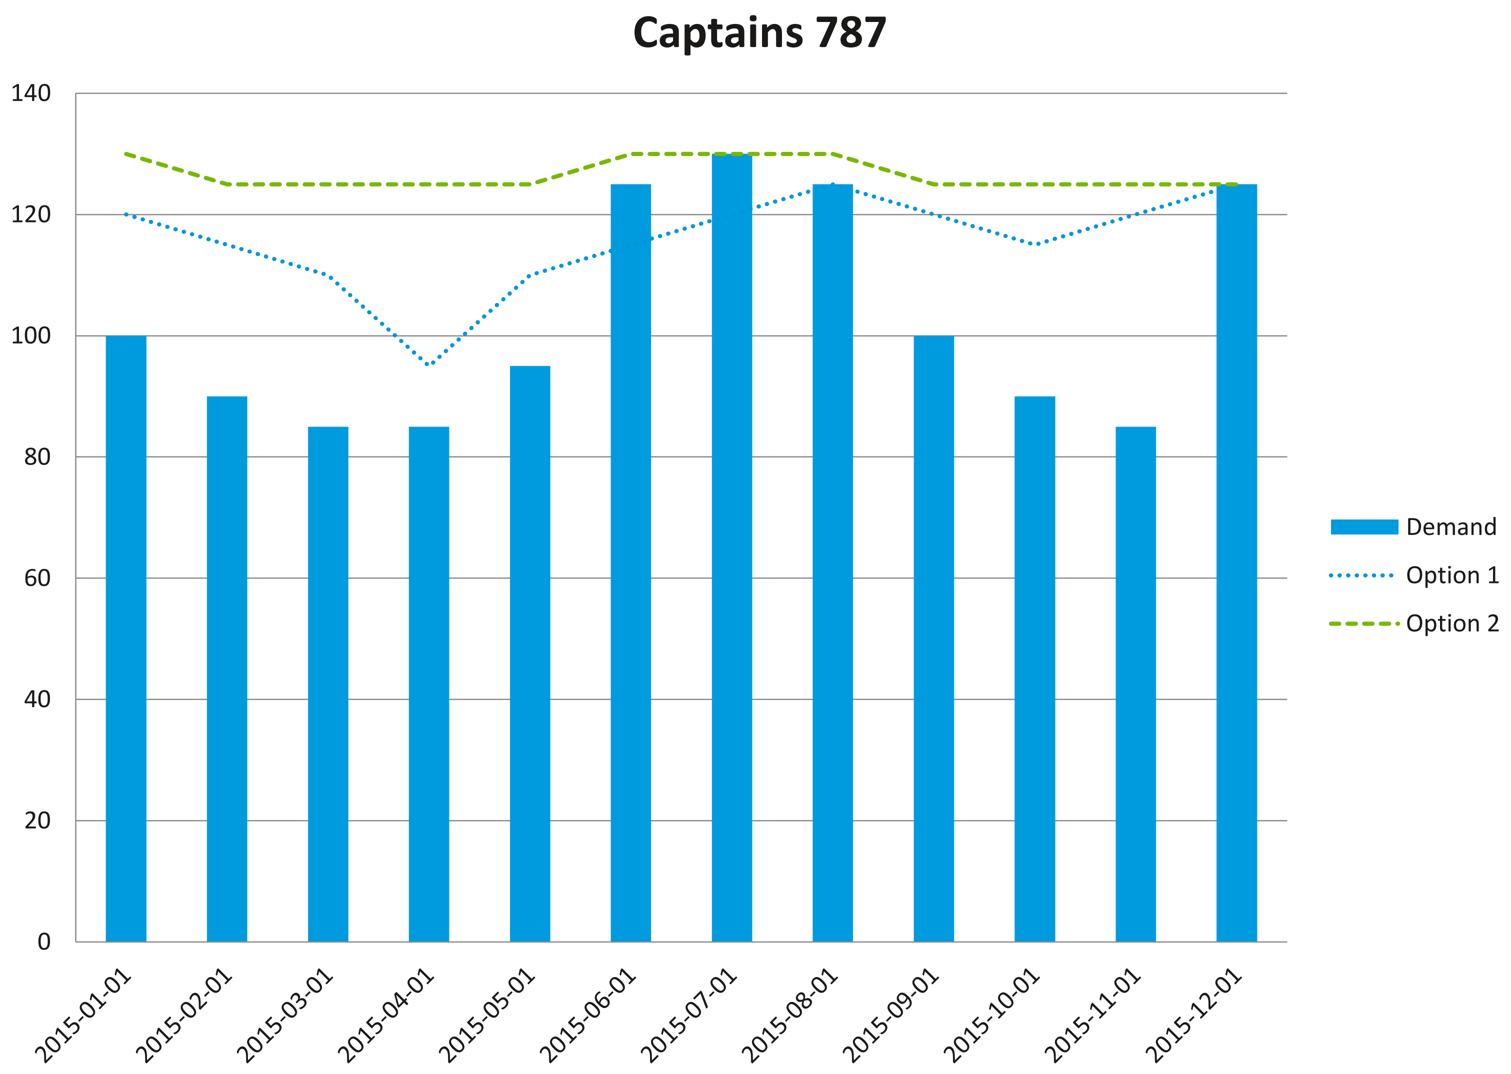

The manpower planning problem

Global Services

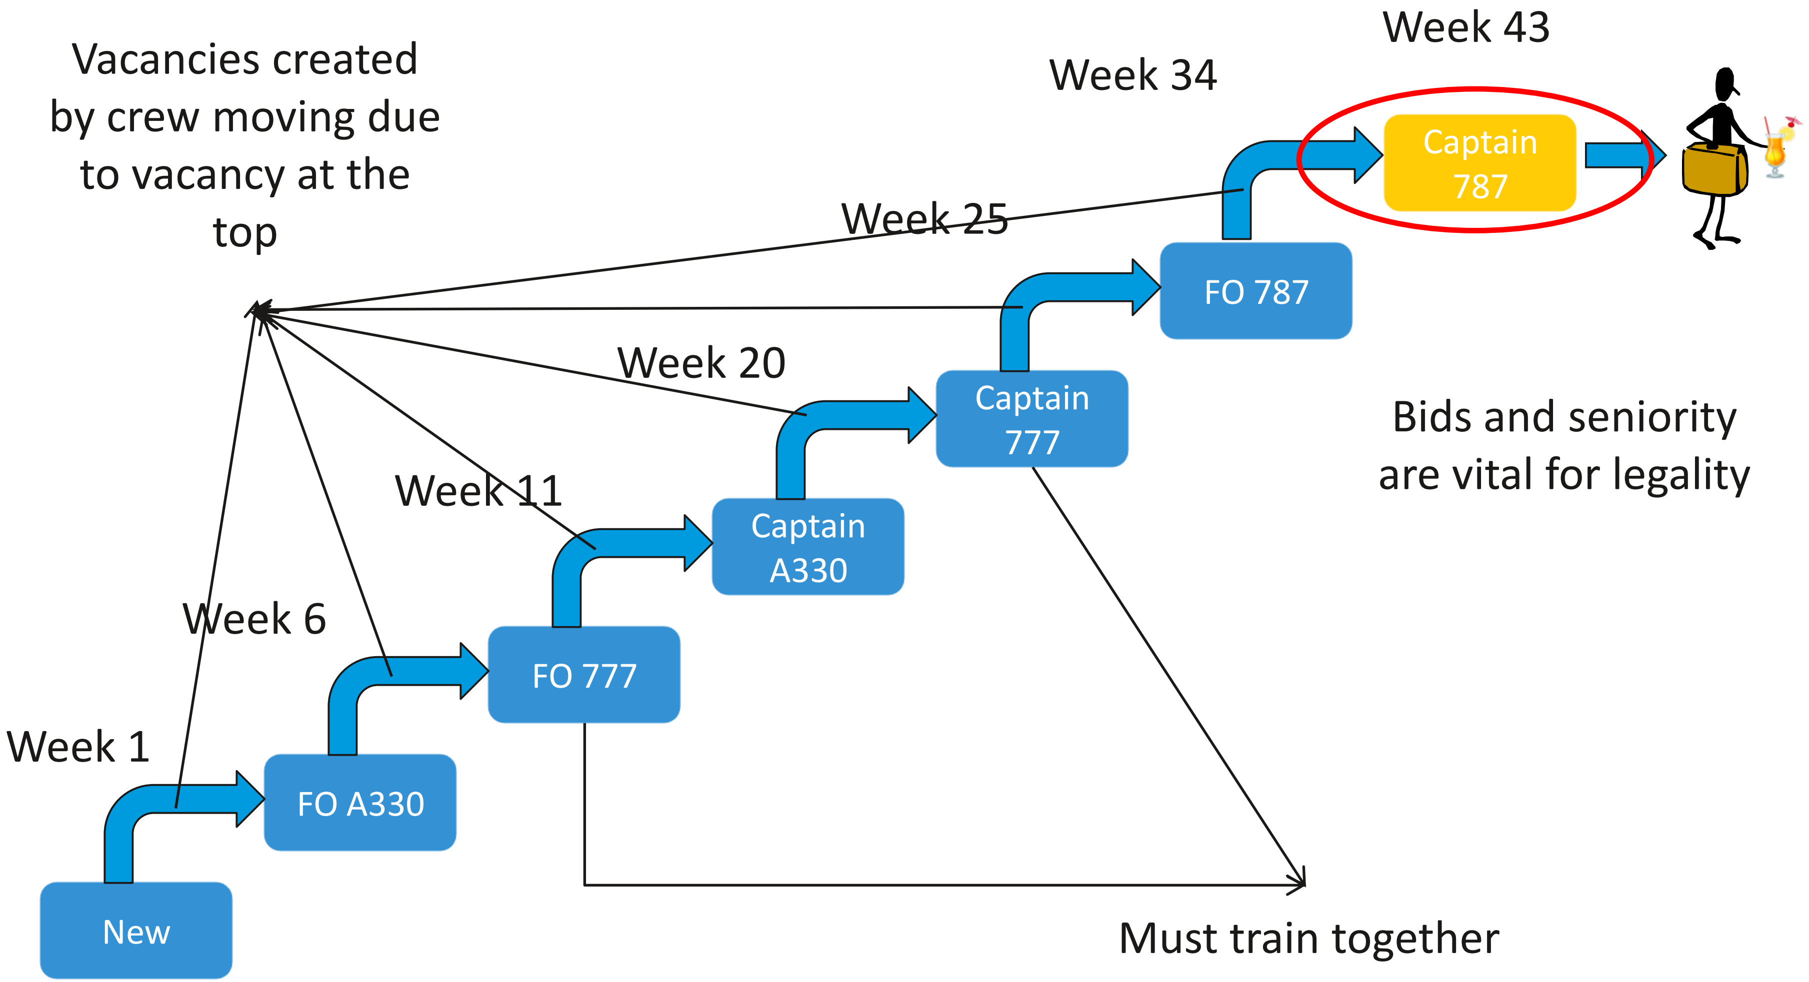

Modeling the manpower planning problem

Global Services

Complications

- Long training periods
- Scarce training resources
    - Simulators, instructors, classrooms
- Strict seniority
- Fleet coordination
- Cost modelling

The problem can be formulated as a MIP!

But it cannot be directly solved in practice, since it is too large!

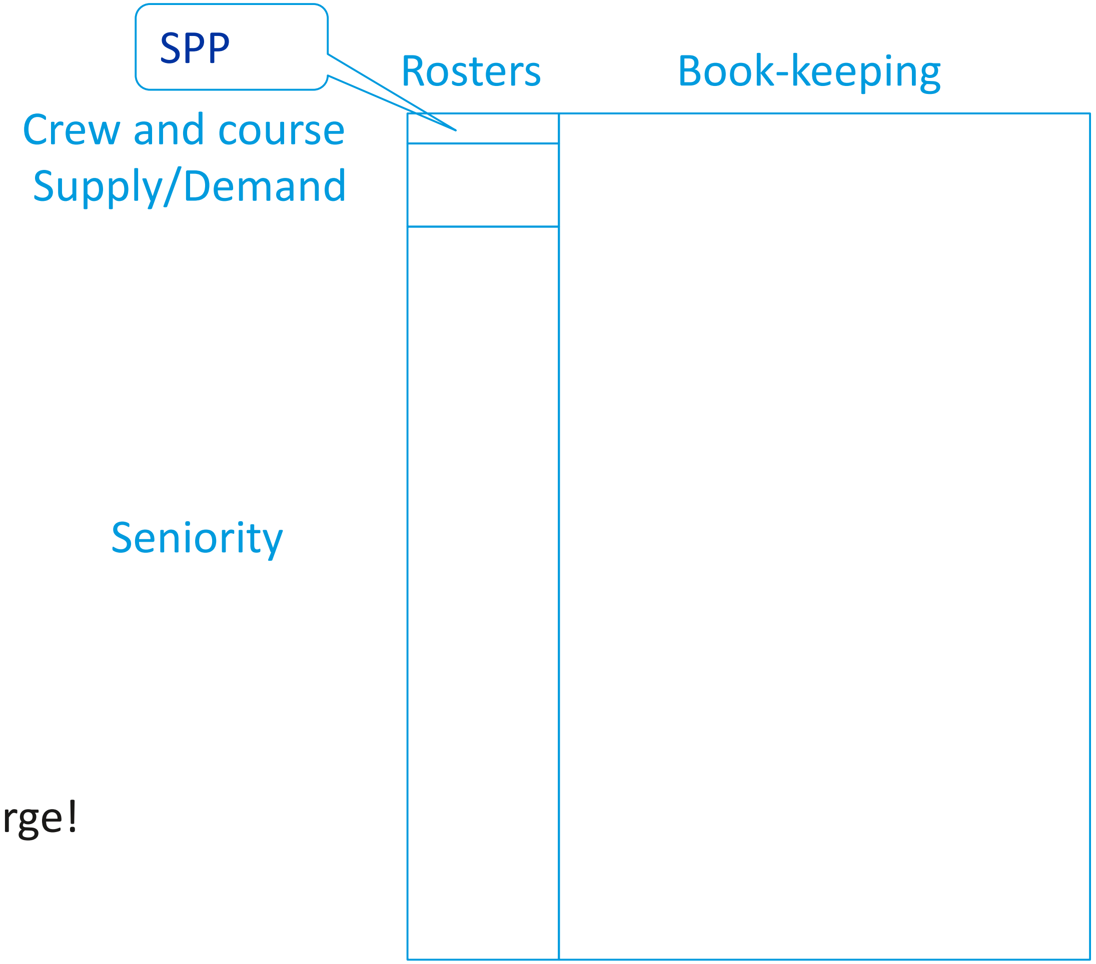

Solving the manpower planning problem

Global Services

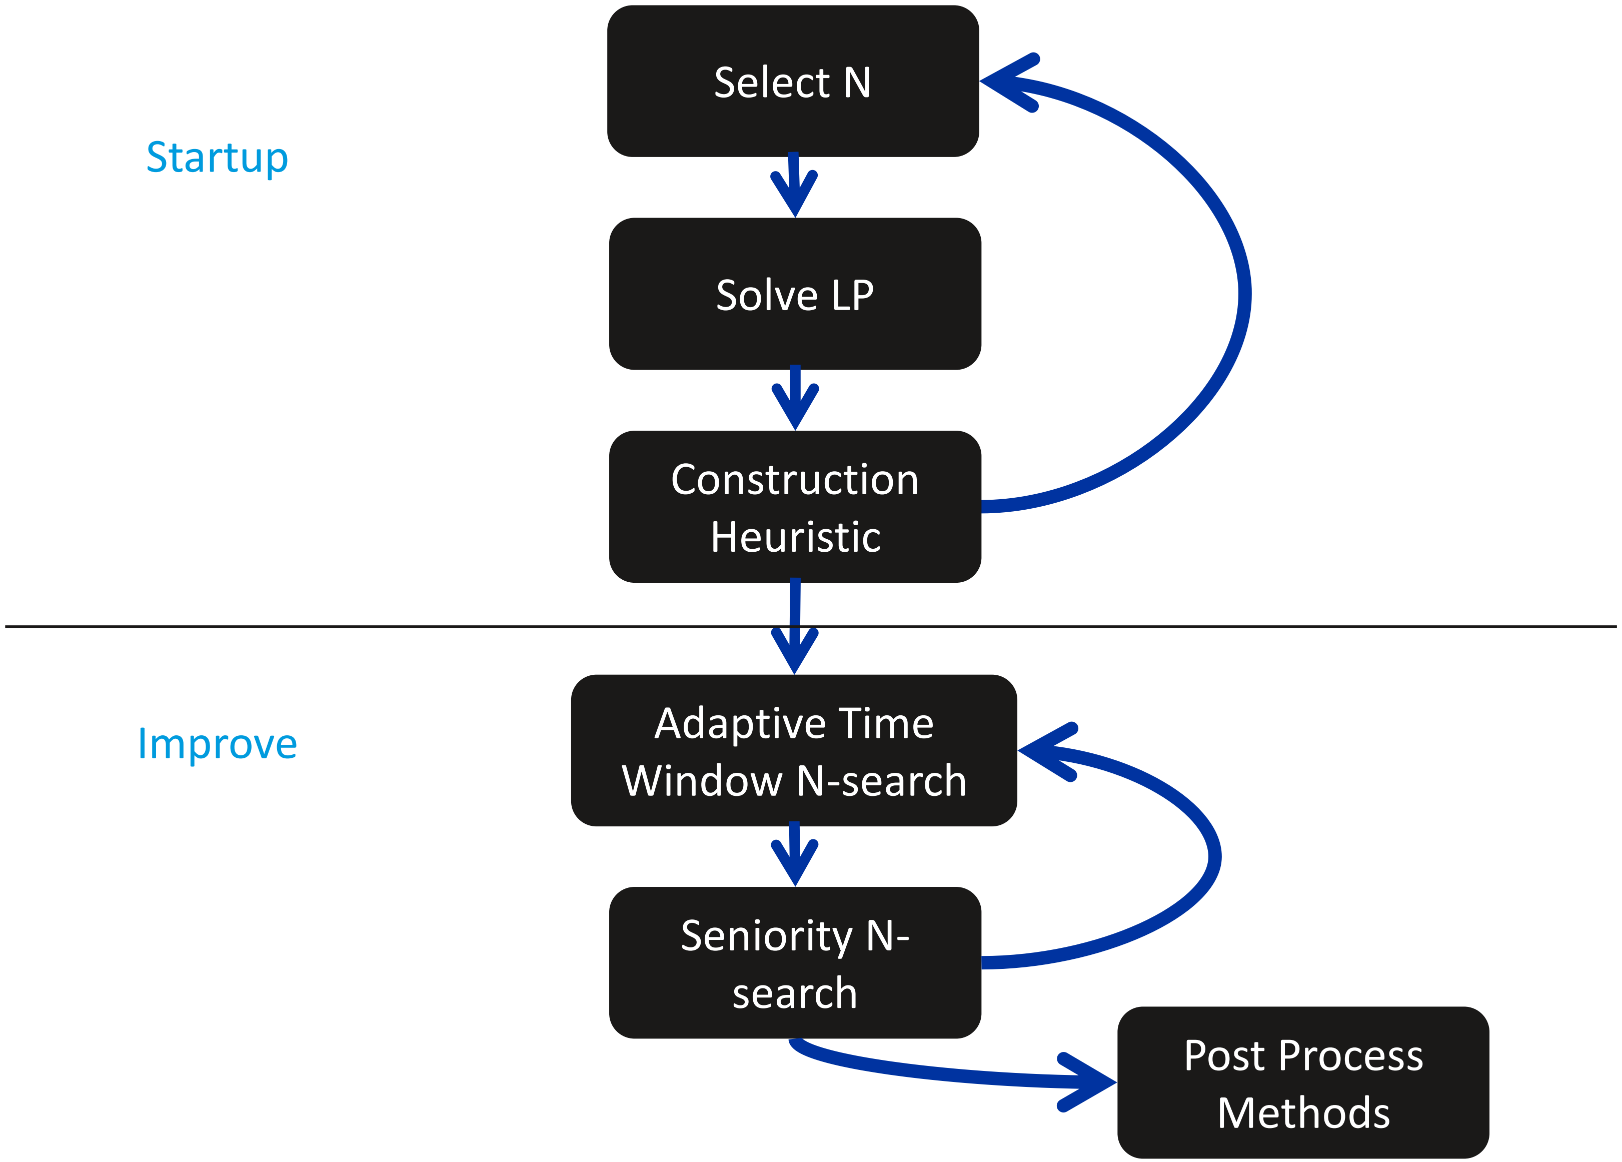

The crew pairing problem

Create anonymous crew ‘work shifts’

In such a way that

- All work shifts (*crew pairings*) start and end at a crew base
- All flights are included in at least one crew pairing
- All safety, union and company rules and regulations are satisfied
- The number of crew pairings for each base is limited by the number of crew stationed at the base
- Crew costs (layover costs, overtime costs, per diem …) are minimized

Crew pairing is the most well-studied airline crew planning problem.

Modeling the crew pairing problem

While the main complication for tail assignment is the combinatorial explosion, the main complication for crew pairing, and crew planning in general, is the number of rules.

Crew pairings are always a few days long at most, and very sparse.

However, they need to satisfy a lot of complicated rules.

Flights can be overcovered (included in more than one pairing), and these overcovers are transformed to so-called *deadheads* (crew flying as passengers)

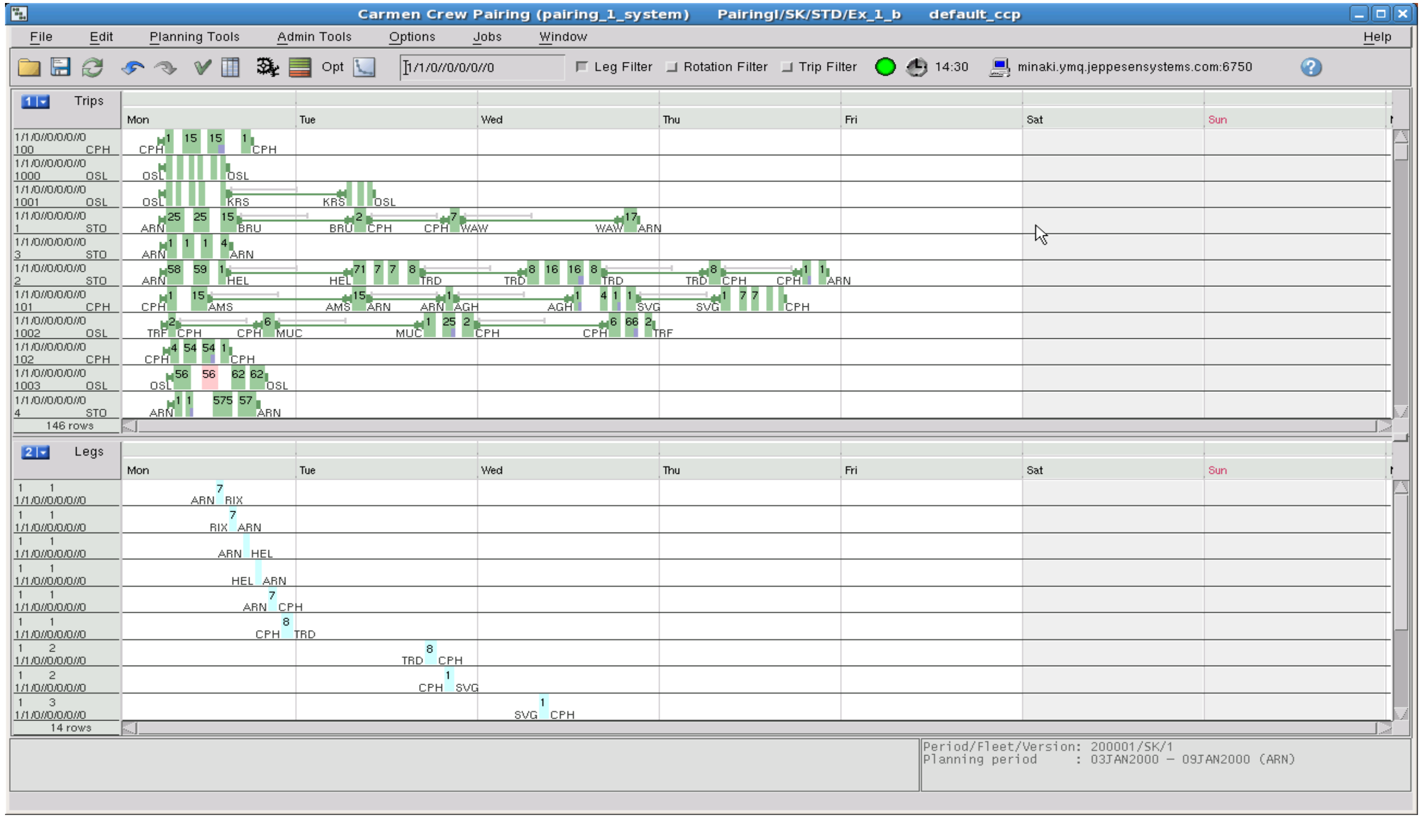

Solving the crew pairing problem

Global Services

The crew pairing problem is solved using column generation, much like tail assignment.

We also use highly specialized integer heuristics to quickly find integer solutions.

Our world-leading crew pairing optimizer is our main product, and the reason we were fonded!

The crew rostering problem

Create personal crew schedules

In such a way that

- All crew pairings are assigned to the right number of crew (one or more)
- Pre-planned vacation, courses etc are respected
- All safety, union and company rules and regulations are satisfied
- Personal requests (*bids*) are taken into consideration
- Some measure of fairness between crew is considered
- Crew costs and crew bid satisfaction is optimized

The crew rostering problem

Global Services

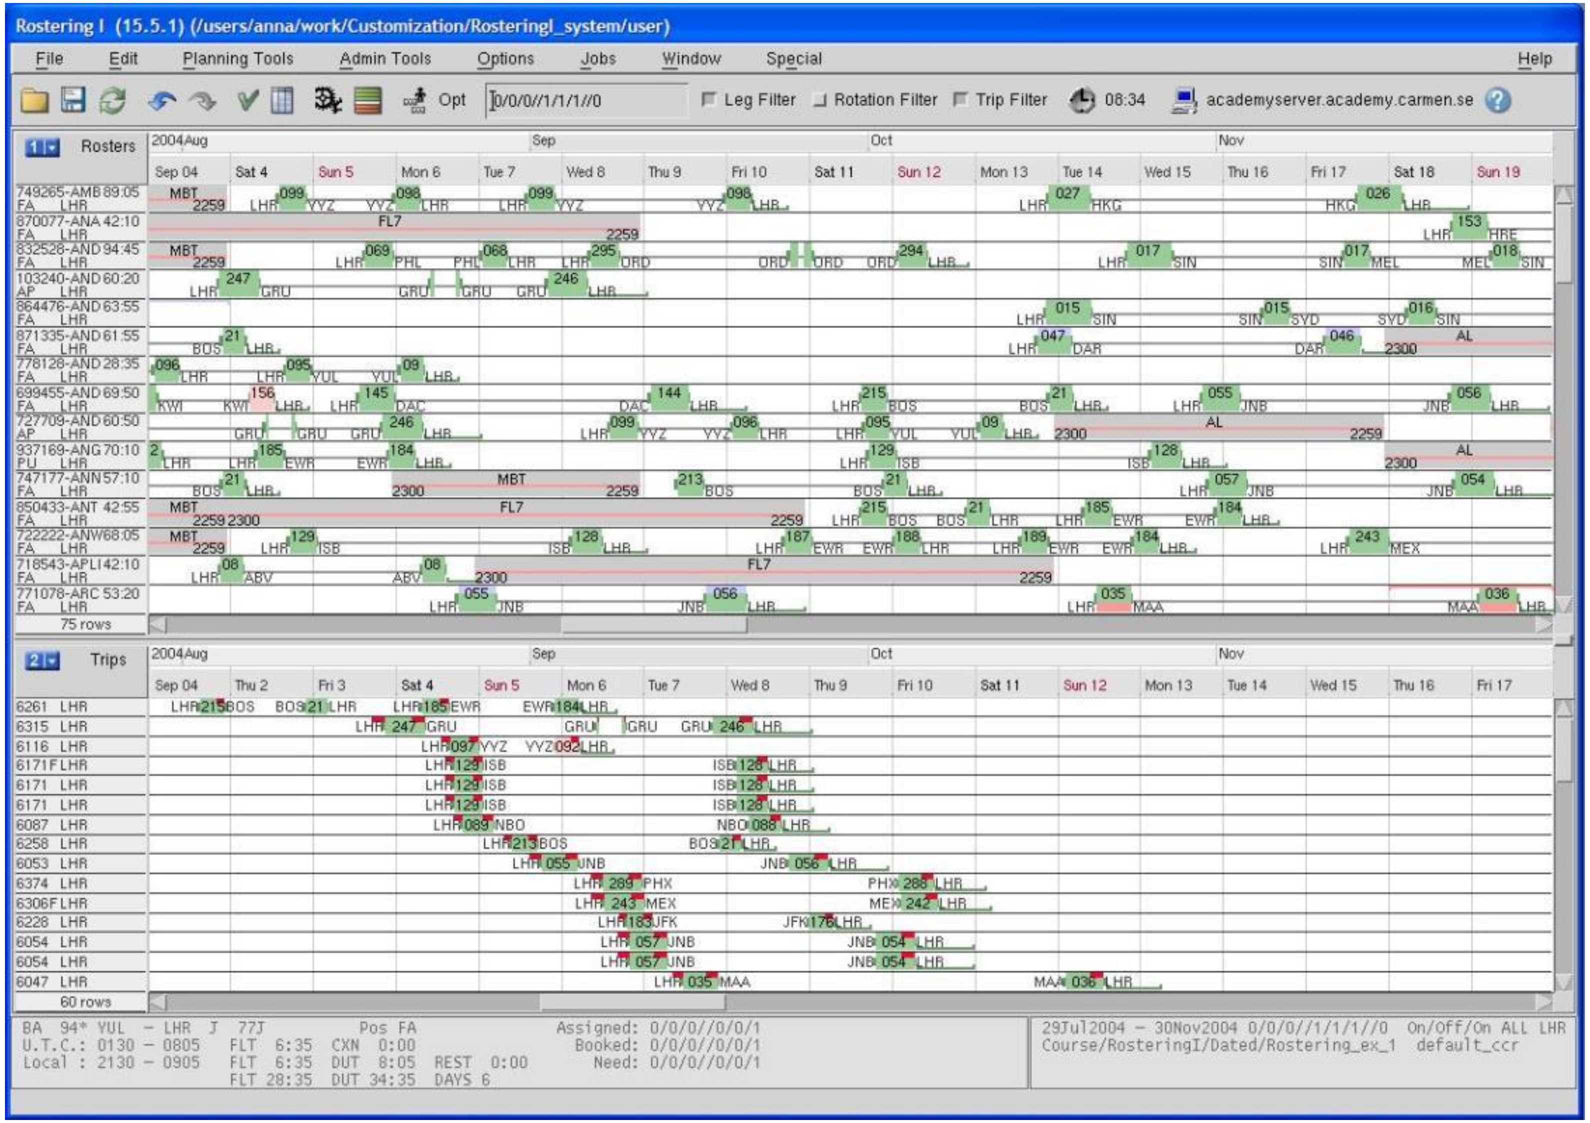

Modeling the crew rostering problem

Global Services

- **Rules**
  - Max block time per month
  - Minimum rest/consecutive work days
  - Passport/visa rules
  - Qualifications

- **Constraints (soft, break at a cost)**
  - Max inexperienced per flight
  - Language constraints
  - Incompatible crew
  - Biddy bids

- **Fairness – fairness each month, or over time?**
  - US: Strict seniority. More senior crew members are ’infinitely more important’ than less senior ones. Fair over a lifetime, extremely unfair each month.
  - Rest of the world: Weighted fairshare. Award each crew member approximately the same amount of what they have bid for. Pretty fair each month.

Solving the crew rostering problem

Global Services

The type of fairness used decides which solution method works best:

Strict seniority

The main difficulty is ensuring that a solution covers all trips, and doesn’t contain any ’seniority inversions’. A seniority inversion means that a less senior crew member gets something which a more senior one bid for. Seniority inversions incur heavy monetery penalties awarded to crew.

Finding seniority inversions is very complicated, and as a result **strict seniority rostering does not belong to the complexity class P –** it is not possible to verify a solution in polynomial time.

MIP based local search:

1. Find a feasible solution using a construction heuristic.
2. Select random subsets of crew, look for and fix seniority inversions.
3. Stop when some convergence criterion is reached.

Solving the crew rostering problem

Global Services

Weighted fairshare

The main difficulty is ensuring that a solution covers all trips and distributes bids fairly over the crew members.

Column generation is the best solution methods for most weighted fairshare rostering problems.

Exceptions are problems with very special rule structure, which ruins the possibilities to efficiently generate new crew rosters. For such methods, very large neighborhood search techniques are used.

Operational problems

Operational problems

Global Services

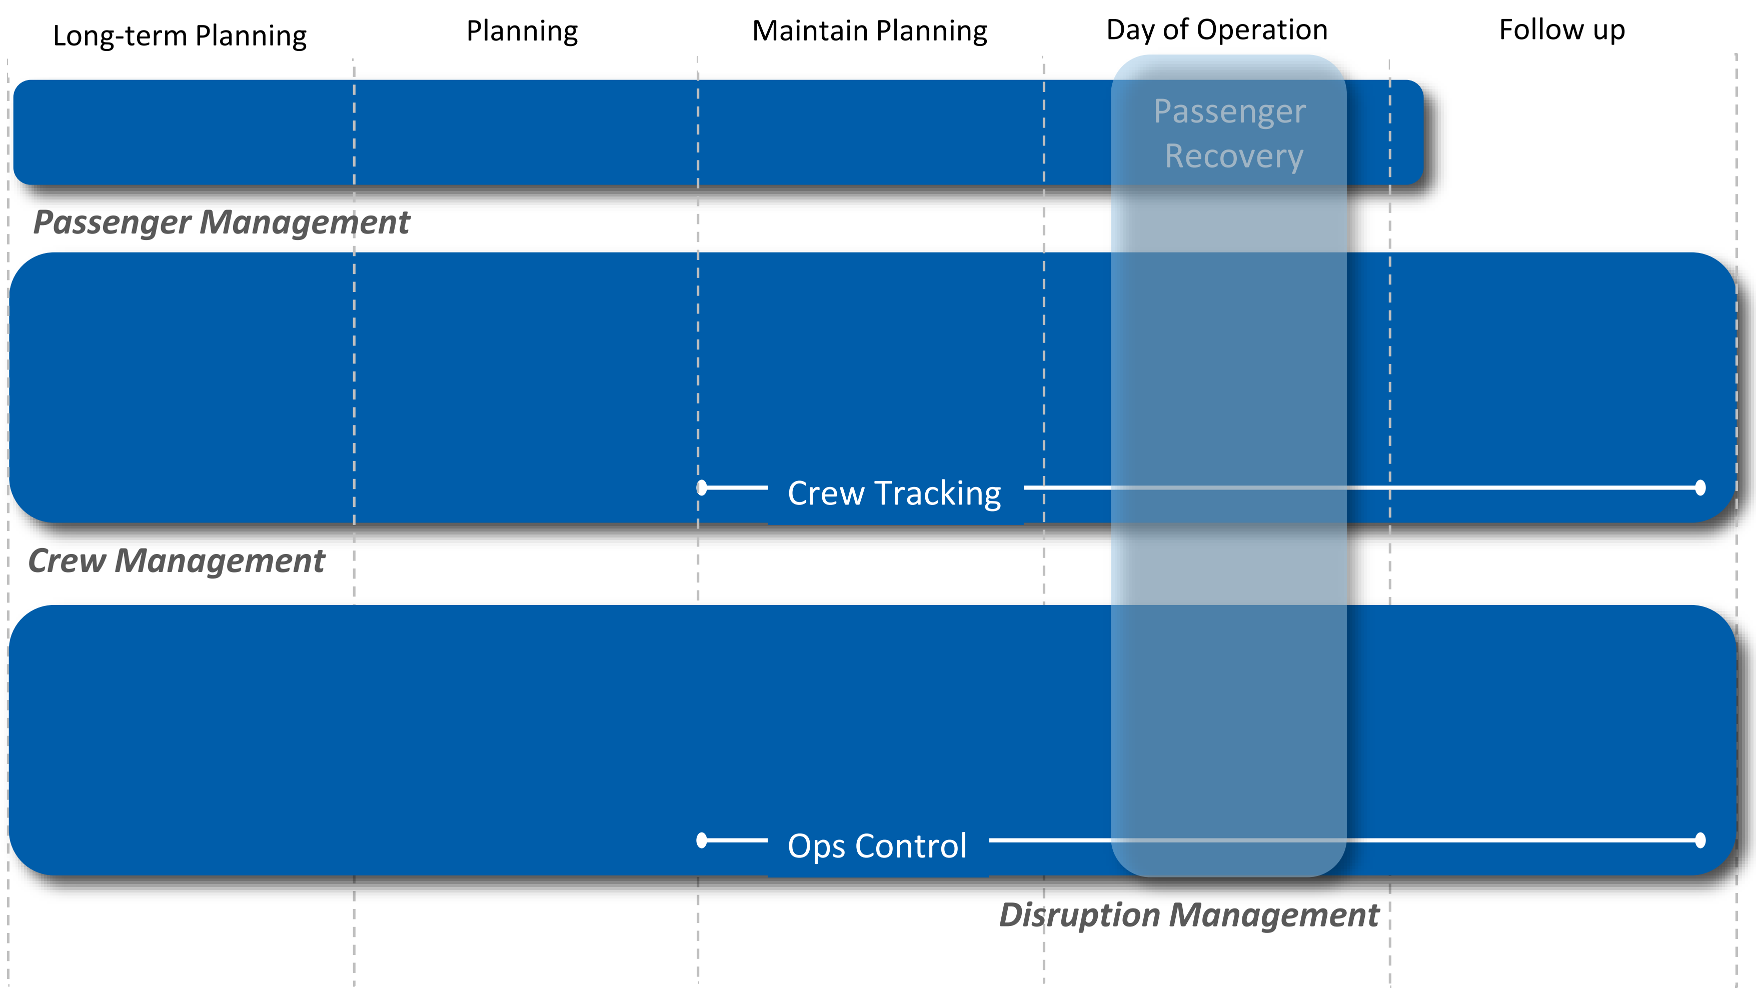

The ops control process

Fix issues with flight and aircraft schedules

In such a way that

- All flights are assigned, if possible
- Passenger effects are minimal
- Crew effects are minimal
- Pre-planned maintenance is respected, unless it's not ☺
- All operational rules are satisfied, unless they're not ☺
- Operation gets back to plan as soon as possible

Ops control is extremely time-critical and driven by external factors – monitor aircraft and airport operations, alert when stuff goes wrong.

Local search is used to quickly produce multiple solution options.

The crew tracking process

Fix issues with personal crew schedules

In such a way that

- All crew pairings are assigned to the right number of crew (one or more)
- Pre-planned vacation, courses etc are respected
- All safety, union and company rules and regulations are satisfied
- Personal requests (*bids*) are taken into consideration
- Some measure of fairness between crew is considered
- Crew costs and crew bid satisfaction is <u>maintained</u>

Crew tracking is time-critical and very scattered – optimization is not enough, you might also need to call extra crew, check how late people will be, change hotel bookings etc… A manual process in most cases.

The flight planning problem

Calculate a fuel-optimal flight plan

Before any flight can take place, a *flight plan* needs to be filed.

A flight plan describes the three dimensional flight of the aircraft from its origin to its destination, and also includes time points.

Considers:

- Up-to-date aircraft specs, including fuel burn
- Up-to-date aerodrome and en-route weather predictions
- Up-to-date airspace restrictions and regulations

Solved using highly specialized shortest path algorithms.

Follow up

Follow up analytics and learning

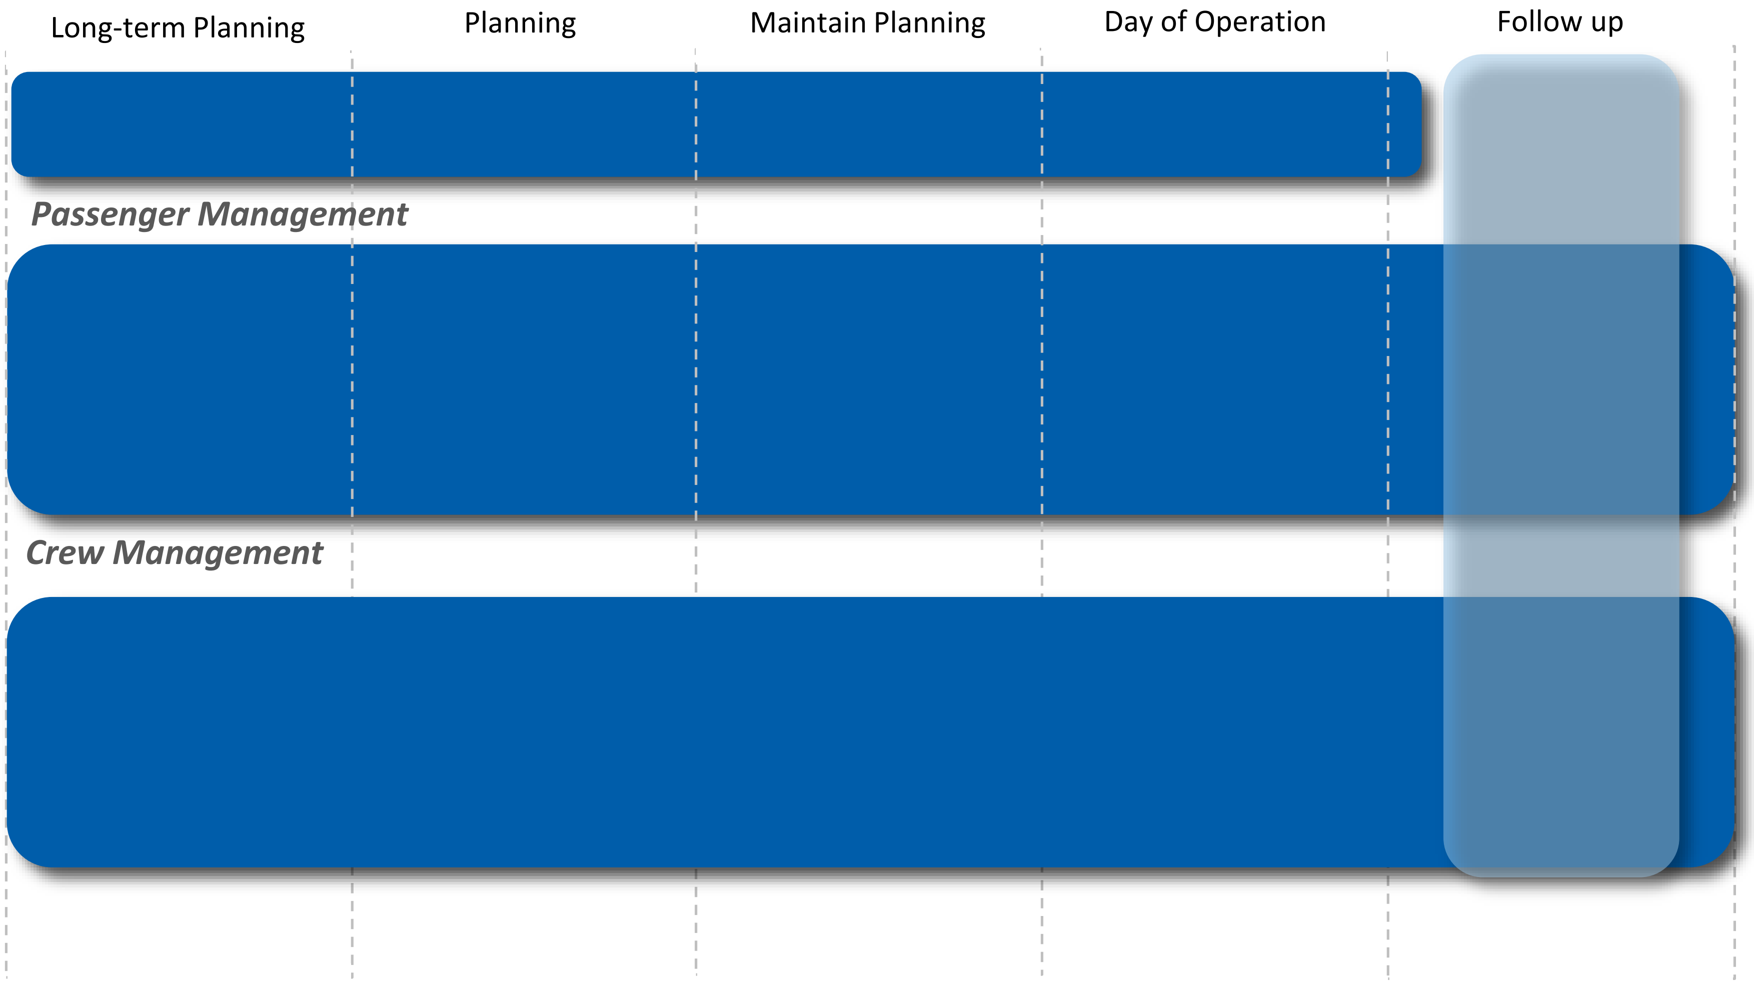

Follow up analytics and learning

Global Services

In the last few years we have increased our efforts in post-operational analytics, follow-up and learning

Since crew and fleet planning are recurring activities, monitoring, learning and adapting is critical.

Examples:

- Monitoring important KPIs
- Predicting flight delays
- Adjusting crew buffer times based on historical data
- Adjusting aircraft and crew schedules to avoid propagated delays
- Learning which operational patterns that work well and which do not work