

Algorithms and Data-Structures III : 1DL481

Lecture 1

Justin Pearson based on slides by Pierre Flener

Outline

- 1 What's happening Today?
- 2 The End of Course AD2
- 3 Combinatorial Optimisation
 - Constraint Problems
 - Solving Technologies
 - Modelling
 - Solving
- 4 Course AD3
 - Contents
 - New Concepts
 - Learning Outcomes
 - Organisation

Today

- From the end of AD2 to AD3. Solving NP problems
- Some other approaches to NP problems (Model + Solve)
- Overview of the course content
- Mechanics

How To Communicate with us.

Please don't use Studium's message service.

- If you have a question about the **lecture material** or **course organisation**, then email me justin.pearson@it.uu.se.
- If you have a question about the **assignments** or **infrastructure**, then contact the assistants at a help session or solution session for an immediate answer.

Short *clarification* questions (that is: *not* about modelling or programming difficulties) that are either emailed (it-ad3@lists.uu.se)

The End of AD2

During AD2 we spent a lot of time trying to find efficient algorithms to problems via Dynamic programs, efficient greedy algorithms that are guaranteed to find a solution.

At the end of AD2 we got to complexity theory. We talk about the class P (polynomial time solvable problems) and the class NP (non-deterministic polynomial time problems).

In this course we are going to look at algorithms to tackle problems in NP efficiently.

Decision Problems

In a *decision problem* we seek a 'yes' / 'no' answer to an existence question. An *instance* of a problem is given by its input data.

Example (Travelling salesperson: Decision TSP)

Given a budget b and a map with n cities, **is** there a route visiting each city exactly once, returning to the starting city, and costing at most b ?

Decision Problems

A decision problem R is:

- *in NP* if a *witness* to a ‘yes’ instance is checkable in time polynomial in the instance size: checking is *in P*;
- *NP-complete* if in NP and there is a *reduction* from each problem Q in NP, polytime converting any instance of Q into a same-answer instance of R .

It is believed that NP-complete problems are *intractable* (or: *hard*), requiring non-polynomial time to solve exactly.

Example

TSP is NP-complete as a witness is checkable in $\mathcal{O}(n)$ time and the NP-complete Hamiltonian-Cycle problem reduces to it.

Satisfaction Problems

In a *satisfaction problem* we seek a witness for a 'yes' answer.

Example (Satisfaction TSP)

Given a budget b and a map with n cities, **find** a route visiting each city exactly once, returning to the starting city, and costing at most b .

Note that finding a witness of a solution could be harder than just deciding if a solution exists or not.

Optimisation Problems

In an *optimisation problem* we seek an optimal witness, according to some *objective function*, for a 'yes' answer.

Example (Optimisation TSP)

Given a map with n cities, **find a cheapest** route visiting each city exactly once and returning to the starting city.

In addition to decision problems that are at least as hard as every NP problem (as every NP problem reduces to them), satisfaction and optimisation problems with NP-complete decision versions are often also said to be *NP-hard*: they are unlikely to be easier than their decision versions.

What Now?

Several courses at Uppsala University teach techniques for addressing NP-hard optimisation and satisfaction problems:

[1TD184](#) Continuous Optimisation (period 2)

[1DL451](#) Modelling for Combinatorial Optimisation (period 2)

[1DL481](#) Algorithms and Data Structures 3 (period 3)

NP-hardness is not where the fun ends, but where it begins!

Example (Optimisation TSP over n cities)

A brute-force algorithm evaluates all $n!$ candidate routes:

- A computer of today evaluates 10^6 routes / second:

n	time
11	40 seconds
14	1 day
18	203 years
20	77k years

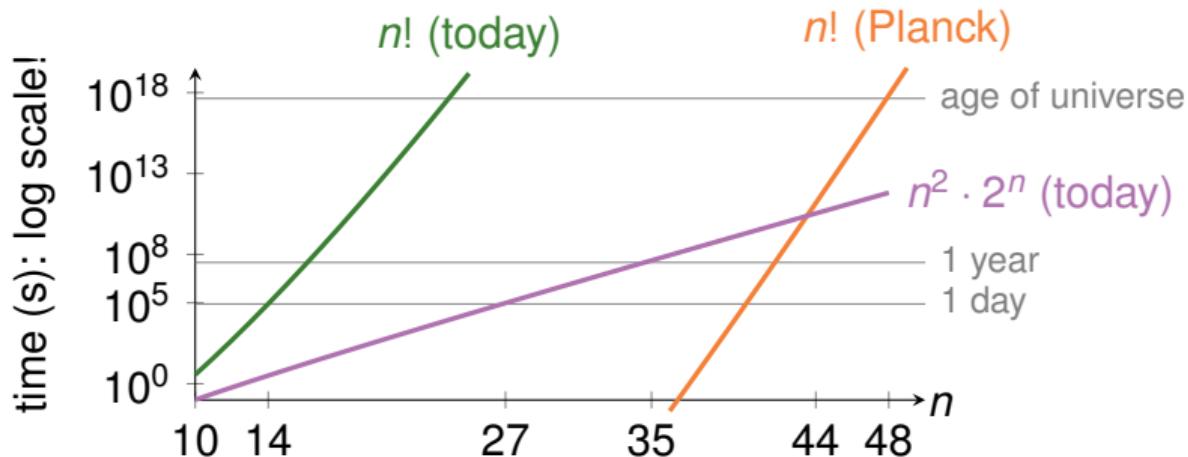
- Planck time is the shortest useful interval: $\approx 5.4 \cdot 10^{-44}$ seconds; a Planck computer would evaluate $1.8 \cdot 10^{43}$ routes / second:

n	time
37	0.7 seconds
41	20 days
48	1.5 · age of universe

The dynamic program by Bellman-Held-Karp “only” takes $\mathcal{O}(n^2 \cdot 2^n)$ time: a computer of today takes a day for $n = 27$, a year for $n = 35$, the age of the universe for $n = 67$, and beats the $\mathcal{O}(n!)$ algo on Planck computer for $n \geq 44$.

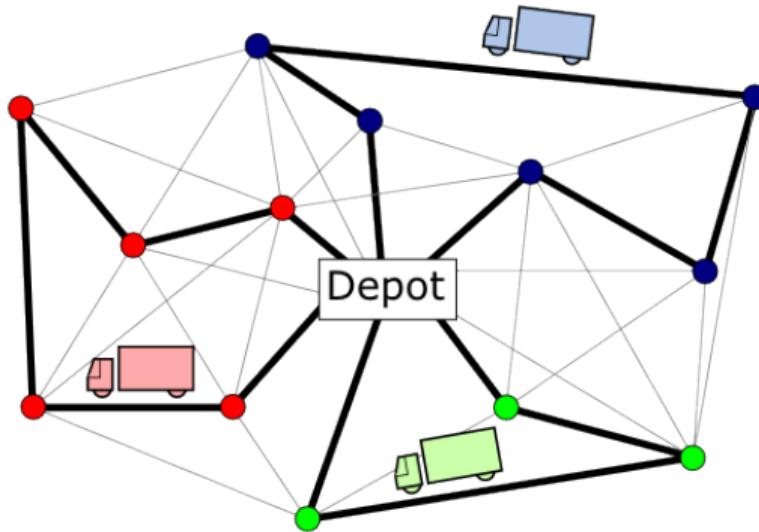
Intelligent Search upon NP-Hardness

Do not give up but try to stay ahead of the curve: there is an instance size until which an **exact** algorithm is fast enough!



Concorde TSP Solver beats *Bellman-Held-Karp* exact algorithm: it uses local search & approximation algos, but sometimes proves exactness of its optima. The largest instance solved exactly, in 136 CPU years in 2006, has $n = 85900$.

Optimisation



Optimisation is a science of **service**: to scientists, to engineers, to artists, and to society.

Example (Agricultural experiment design)

	plot1	plot2	plot3	plot4	plot5	plot6	plot7
barley							
corn							
millet							
oats							
rye							
spelt							
wheat							

Constraints to be satisfied:

- 1 Equal growth load: Every plot grows 3 grains.
- 2 Equal sample size: Every grain is grown in 3 plots.
- 3 Balance: Every grain pair is grown in 1 common plot.

Instance: 7 plots, 7 grains, 3 grains/plot, 3 plots/grain, balance 1.

Example (Agricultural experiment design)

	plot1	plot2	plot3	plot4	plot5	plot6	plot7
barley	✓	✓	✓	—	—	—	—
corn	✓	—	—	✓	✓	—	—
millet	✓	—	—	—	—	✓	✓
oats	—	✓	—	✓	—	✓	—
rye	—	✓	—	—	✓	—	✓
spelt	—	—	✓	✓	—	—	✓
wheat	—	—	✓	—	✓	✓	—

Constraints to be satisfied:

- 1 Equal growth load: Every plot grows 3 grains.
- 2 Equal sample size: Every grain is grown in 3 plots.
- 3 Balance: Every grain pair is grown in 1 common plot.

Instance: 7 plots, 7 grains, 3 grains/plot, 3 plots/grain, balance 1.

Example (Doctor rostering)

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Doctor A							
Doctor B							
Doctor C							
Doctor D							
Doctor E							

Constraints to be satisfied:

- ① #on-call doctors / day = 1
- ② #operating doctors / weekday ≤ 2
- ③ #operating doctors / week ≥ 7
- ④ #appointed doctors / week ≥ 4
- ⑤ day off after operation day
- ⑥ ...

Objective function to be minimised: Cost: ...

Example (Doctor rostering)

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Doctor A	call	none	oper	none	oper	none	none
Doctor B	<i>appt</i>	call	none	oper	none	none	call
Doctor C	oper	none	call	<i>appt</i>	<i>appt</i>	call	none
Doctor D	<i>appt</i>	oper	none	call	oper	none	none
Doctor E	oper	none	oper	none	call	none	none

Constraints to be satisfied:

- ① #on-call doctors / day = 1
- ② #operating doctors / weekday ≤ 2
- ③ #operating doctors / week ≥ 7
- ④ #appointed doctors / week ≥ 4
- ⑤ day off after operation day
- ⑥ ...



Objective function to be minimised: Cost: ...

Example (Vehicle routing: parcel delivery)

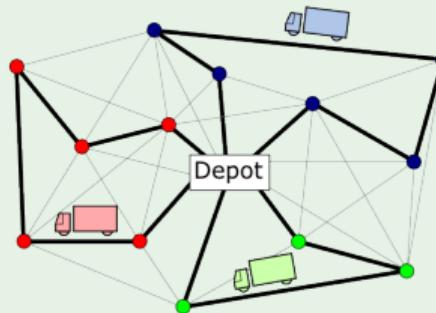
Given a depot with parcels for clients and a vehicle fleet, **find** which vehicle visits which client when.

Constraints to be satisfied:

- 1 All parcels are delivered on time.
- 2 No vehicle is overloaded.
- 3 Driver regulations are respected.
- 4 ...

Objective function to be minimised:

- Cost: the total fuel consumption and driver salary.



Application Areas

School timetabling

	Monday	Tuesday	Wednesday	Thursday	Friday
9.00	MF2202 Ordinary Differential Equations I FH1		LABC 82072 Computer Graphics (B) FH1	MF2203 Numerical Analysis I FH1, FH2	
10.00	KMT2202 Ordinary Differential Equations I FH1, FH2, FH3		LABC 82072 Computer Graphics (B) FH1	KMT2202 Ordinary Differential Equations I FH1, FH2, FH3 KMT2202 Numerical Analysis I FH1	KMT2202 Ordinary Differential Equations I FH1
11.00	CAS012 Algorithms and Data Structures I FH1		MF2202 Numerical Analysis I FH1	MF2202 Ordinary Differential Equations I FH1, FH2	
12.00	MF2212 Further Linear Algebra FH1, FH2, FH3	MF2202 Numerical Analysis I FH1, FH2	CAS012 Computer Graphics I FH1	MF2212 Further Linear Algebra FH1, FH2	MF2212 Further Linear Algebra FH1, FH2
13.00			MF2202 Numerical Analysis I FH1, FH2	MF2212 Further Linear Algebra FH1, FH2	KMT2202 Ordinary Differential Equations I FH1, FH2
14.00				MF2212 Further Linear Algebra FH1, FH2	KMT2202 Ordinary Differential Equations I FH1, FH2
15.00	CAS012 Computer Graphics I FH1		MF2202 Numerical Analysis I FH1, FH2		
16.00		CAS012 Computer Graphics I FH1			
17.00		CAS012 Algorithms and Data Structures I FH1			

Security: SQL injection?



Sports tournament design

svensk handball



Container packing



Applications in Programming and Testing

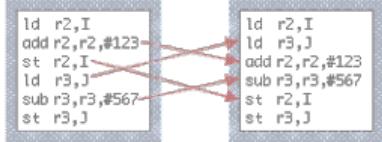
Robot programming



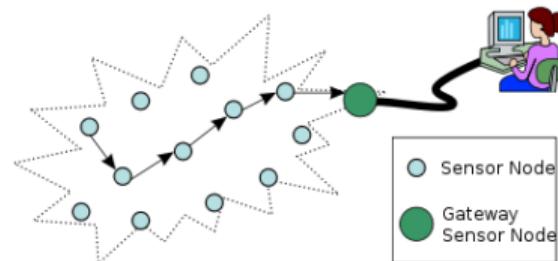
Compiler design

COMPILERS
FOR INSTRUCTION SCHEDULING

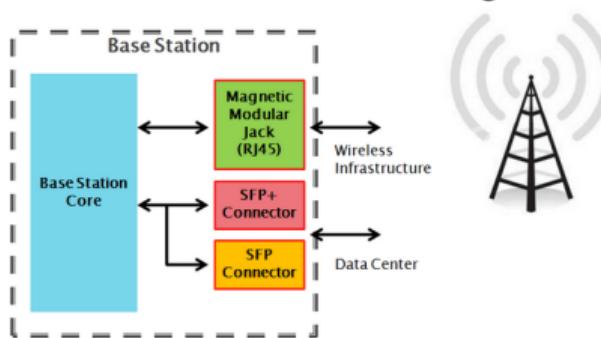
C Compiler
C++ Compiler



Sensor-net configuration

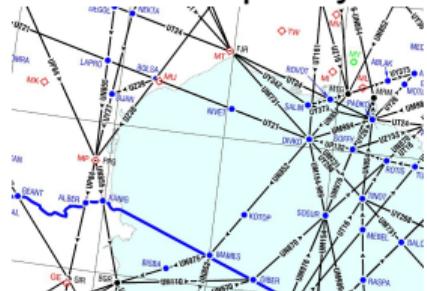


Base-station testing



Applications in Air Traffic Management

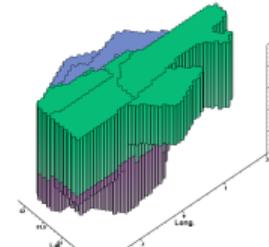
Demand vs capacity



Contingency planning

Flow	Time Span	Hourly Rate
From: Arlanda	00:00 – 09:00	3
	09:00 – 18:00	5
	18:00 – 24:00	2
From: Arlanda	00:00 – 12:00	4
	12:00 – 24:00	3
...

Airspace sectorisation

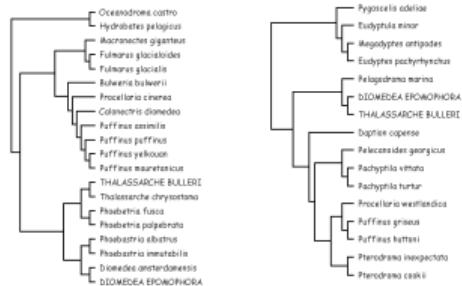


Workload balancing



Applications in Biology and Medicine

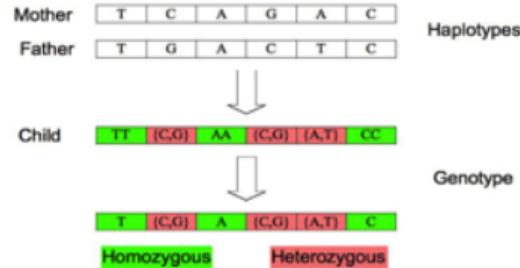
Phylogenetic supertree



Medical image analysis



Haplotype inference



Doctor rostering



Definitions

In a *constraint problem*, values have to be **found** for all the unknowns, called *variables* (in the mathematical sense; also called *decision variables*) and ranging over **given** sets, called *domains*, so that:

- All the given *constraints* on the decision variables are satisfied.
- Optionally: A given *objective function* on the decision variables has an optimal value: either a minimal cost or a maximal profit.

A *candidate solution* to a constraint problem maps each decision variable to a value within its domain; it is:

- *feasible* if all the constraints are satisfied;
- *optimal* if the objective function takes an optimal value.

The *search space* consists of all candidate solutions. A *solution* to a *satisfaction problem* is feasible. An *optimal solution* to an *optimisation problem* is feasible and optimal.

Search spaces are often larger than the universe!



Many important real-life problems are NP-hard or worse: their real-life instances can only be solved exactly and fast enough by **intelligent** search, unless $P = NP$. **NP-hardness is not where the fun ends, but where it begins!**

A *solving technology* offers languages, methods, and tools for:

what: Modelling constraint problems in a declarative language.

and / or

how: Solving constraint problems intelligently:

- Search: Explore the space of candidate solutions.
- Inference: Reduce the space of candidate solutions.
- Relaxation: Exploit solutions to easier problems.

A *solver* is an off-the-shelf program that takes any model and data as input and tries to solve that problem instance.

Combinatorial (= discrete) optimisation covers satisfaction and optimisation problems for variables ranging over *discrete* sets: *combinatorial problems*.

Example (Agricultural experiment design, AED)

	plot1	plot2	plot3	plot4	plot5	plot6	plot7
barley	✓	✓	✓	—	—	—	—
corn	✓	—	—	✓	✓	—	—
millet	✓	—	—	—	—	✓	✓
oats	—	✓	—	✓	—	✓	—
rye	—	✓	—	—	✓	—	✓
spelt	—	—	✓	✓	—	—	✓
wheat	—	—	✓	—	✓	✓	—

Constraints to be satisfied:

- 1 Equal growth load: Every plot grows 3 grains.
- 2 Equal sample size: Every grain is grown in 3 plots.
- 3 Balance: Every grain pair is grown in 1 common plot.

Instance: 7 plots, 7 grains, 3 grains/plot, 3 plots/grain, balance 1.

General term: *balanced incomplete block design (BIBD)*.

Example (Agricultural experiment design, AED)

	plot1	plot2	plot3	plot4	plot5	plot6	plot7
barley	1	1	1	0	0	0	0
corn	1	0	0	1	1	0	0
millet	1	0	0	0	0	1	1
oats	0	1	0	1	0	1	0
rye	0	1	0	0	1	0	1
spelt	0	0	1	1	0	0	1
wheat	0	0	1	0	1	1	0

Constraints to be satisfied:

- 1 Equal growth load: Every plot grows 3 grains.
- 2 Equal sample size: Every grain is grown in 3 plots.
- 3 Balance: Every grain pair is grown in 1 common plot.

Instance: 7 plots, 7 grains, 3 grains/plot, 3 plots/grain, balance 1.

General term: *balanced incomplete block design (BIBD)*.

In a BIBD, the plots are called *blocks* and the grains are called *varieties*:

Example (BIBD *integer* model: $\checkmark \rightsquigarrow 1$ and $- \rightsquigarrow 0$)

```
-3 enum Varieties; enum Blocks;
-2 int: blockSize; int: sampleSize; int: balance;
-1 array[Varieties,Blocks] of var 0..1: BIBD; % BIBD[v,b]=1 iff v is in b
0 solve satisfy;
1 constraint forall(b in Blocks) (blockSize = count(BIBD[..,b], 1));
2 constraint forall(v in Varieties) (sampleSize = count(BIBD[v,..], 1));
3 constraint forall(v, w in Varieties where v < w) (balance =
  count([BIBD[v,b]+BIBD[w,b] | b in Blocks], 2));
```

Example (Instance data for our AED)

```
-3 Varieties = {barley,...,wheat}; Blocks = {plot1,...,plot7};
-2 blockSize = 3; sampleSize = 3; balance = 1;
```

Reconsider the model fragment:

```
2 constraint forall(v in Varieties) (sampleSize = count(BIBD[v,...], 1));
```

This constraint is declarative (and by the way not within linear algebra), so read it using only the verb “to be” or synonyms thereof:

for all varieties v, the count of occurrences of 1 in row v of BIBD must equal sampleSize

The constraint is not procedural:

for all varieties v, we first count the occurrences of 1 in row v and then check if that count equals sampleSize

The latter reading is appropriate for solution checking, but solution finding performs no such procedural counting.

Example (Idea for another BIBD model)

barley	{plot1, plot2, plot3}	}
corn	{plot1, plot4, plot5}	}
millet	{plot1, plot6, plot7}	}
oats	{plot2, plot4, plot6}	}
rye	{plot2, plot5, plot7}	}
spelt	{plot3, plot4, plot7}	}
wheat	{plot3, plot5, plot6}	}

Constraints to be satisfied:

- 1 Equal growth load: Every plot grows 3 grains.
- 2 Equal sample size: Every grain is grown in 3 plots.
- 3 Balance: Every grain pair is grown in 1 common plot.

Example (BIBD *set* model: a block *set* per variety)

```
-3 enum Varieties; enum Blocks;
-2 int: blockSize; int: sampleSize; int: balance;
-1 array[Varieties] of var set of Blocks: BIBD; % BIBD[v] = blocks for v
0 solve satisfy;
1 constraint forall(b in Blocks)                                (blockSize   =
  sum(v in Varieties) (b in BIBD[v]));
2 constraint forall(v in Varieties)                                (sampleSize =
  card(BIBD[v]));
3 constraint forall(v, w in Varieties where v < w)                (balance     =
  card(BIBD[v] intersect BIBD[w]));
```

Example (Instance data for our AED)

```
-3 Varieties = {barley, ..., wheat}; Blocks = {plot1, ..., plot7};
-2 blockSize = 3; sampleSize = 3; balance = 1;
```

Example (Doctor rostering)

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Doctor A	call	none	oper	none	oper	none	none
Doctor B	<i>appt</i>	call	none	oper	none	none	call
Doctor C	oper	none	call	<i>appt</i>	<i>appt</i>	call	none
Doctor D	<i>appt</i>	oper	none	call	oper	none	none
Doctor E	oper	none	oper	none	call	none	none

Constraints to be satisfied:

- ① #on-call doctors / day = 1
- ② #operating doctors / weekday ≤ 2
- ③ #operating doctors / week ≥ 7
- ④ #appointed doctors / week ≥ 4
- ⑤ day off after operation day
- ⑥ ...



Objective function to be minimised: Cost: ...

Example (Doctor rostering)

```
-5 set of int: Days;  % d mod 7 = 1 iff d is a Monday
-4 enum Doctors;
-3 enum ShiftTypes = {appt, call, oper, none};
-2 % Roster[i,j] = shift type of Dr i on day j:
-1 array[Doctors,Days] of var ShiftTypes: Roster;
0 solve minimize ...;  % plug in an objective function
1 constraint forall(d in Days) (count(Roster[..,d],call) = 1);
2 constraint forall(d in Days where d mod 7 in 1..5)
    (count(Roster[..,d],oper) <= 2);
3 constraint count(Roster,oper) >= 7;
4 constraint count(Roster,appt) >= 4;
5 constraint forall(d in Doctors)
    (regular(Roster[d,..],"((oper none) | appt | call | none)*"));
6 ...  % other constraints
```

Example (Instance data for our small hospital unit)

```
-5 Days = 1..7;
-4 Doctors = {Dr_A, Dr_B, Dr_C, Dr_D, Dr_E};
```

Example (Sudoku)

8								
	3	6						
7		9	2					
5			7					
		4	5	7				
		1			3			
	1				6	8		
	8	5			1			
9				4				

8	1	2	7	5	3	6	4	9
9	4	3	6	8	2	1	7	5
6	7	5	4	9	1	2	8	3
1	5	4	2	3	7	8	9	6
3	6	9	8	4	5	7	2	1
2	8	7	1	6	9	5	3	4
5	2	1	9	7	4	3	6	8
4	3	8	5	2	6	9	1	7
7	9	6	3	1	8	4	5	2

```
-2 array[1..9,1..9] of var 1..9: Sudoku;
-1 ... % load the hints
0 solve satisfy;
1 constraint forall(row in 1..9) (all_different(Sudoku[row,...]));
2 constraint forall(col in 1..9) (all_different(Sudoku[..,col]));
3 constraint forall(i,j in {0,3,6}) (all_different(Sudoku[i+1..i+3,j+1..j+3]));
```

Modelling Languages

The following fully declarative modelling languages are powerful enough to encode NP-hard problems:

- *Mixed integer programming (MIP)*: satisfy a set of linear equalities ($=$) and inequalities ($<$, \leq , \geq , $>$), but not disequalities (\neq), over real-number decision variables and integer decision variables weighted by real-number constants, such that a linear objective function is optimised.
- *Boolean satisfiability solving (SAT)*: satisfy a set of disjunctions of possibly negated Boolean decision variables.
- SAT modulo theories (SMT) and constraint programming (CP) do not have such small standardised low-level modelling languages, but enable the higher level of the previous sample models.

In course 1DL451: Modelling for Combinatorial Optimisation, we use such higher-level models in order to drive CP, MIP, SAT, SMT, ... solvers.

Examples (Solving technologies)

With general-purpose solvers, taking model and data as input:

- (Mixed) integer linear programming (IP and MIP)
- Boolean satisfiability (SAT)
- SAT (resp. optimisation) modulo theories (SMT and OMT)
- Constraint programming (CP)
- ...
- Hybrid technologies (LCG = CP + SAT, ...)

Methodologies, *usually without* modelling and solvers:

- Dynamic programming (DP)
- Greedy algorithms
- Approximation algorithms
- Stochastic local search (SLS)
- ...

Examples (Solving technologies)

With general-purpose solvers, taking model and data as input:

- (Mixed) integer linear programming (IP and MIP) in AD3
- Boolean satisfiability (SAT) in AD3
- SAT (resp. optimisation) modulo theories (SMT and OMT) SMT in AD3
- Constraint programming (CP) via 1DL705
- ...
- Hybrid technologies (LCG = CP + SAT, ...)

Methodologies, *usually without* modelling and solvers:

- Dynamic programming (DP) in 1DL231: AD2
- Greedy algorithms in 1DL231: AD2
- Approximation algorithms in AD3
- Stochastic local search (SLS) in AD3
- ...

Solvers

- *Black-box solvers* (for SAT, SMT, OMT, IP, MIP, ...) have general-purpose search + inference + relaxation that is difficult to influence by the modeller.

AD3

- *Glass-box solvers* (for CP, LCG, ...) have general-purpose search + inference + relaxation that is easy to influence, if desired, by the modeller.

via 1DL705

- *Special-purpose solvers* (for TSP, ...) exist for pure problems (that is: problems without side constraints).

Modelling is an Art

There are good and bad models for each constraint problem:

AD3 and 1DL451: Modelling

- Different models of a problem may take different time on the same solver for the same instance.
- Different models of a problem may scale differently on the same solver for instances of growing size.
- Different solvers may take different time on the same model for the same instance.

Good modellers are worth their weight in gold!

Use solvers: based on decades of cutting-edge research, they are very hard to beat on exact solving.

Common Thread: Coping with NP-Hardness

- ① Mixed integer programming (MIP)
- ② Stochastic local search (SLS)
- ③ Amortised analysis (CLRS4: Chapter 16)
- ④ Probabilistic analysis (Chapter 5)
- ⑤ Randomised algorithms: (Chapter 5) universal hashing, ... (Section 11.3.4)
- ⑥ Proving NP-completeness by reduction (Chapter 34)
- ⑦ Approximation algorithms (Chapter 35)
- ⑧ Boolean satisfiability (SAT)
- ⑨ SAT modulo theories (SMT)

CLRS4 Textbook:

Introduction to Algorithms (4th edition) (errata).

T. H. Cormen, Ch. E. Leiserson, R. L. Rivest, and C. Stein.

The MIT Press, 2022.

In a *probabilistic algorithm analysis*, we use probability theory: knowing or assuming the distribution of the inputs, we compute the *average-case time* (as opposed to the worst-case time) of a deterministic algorithm.

Example

The brute-force string matching algorithm for finding all occurrences of a pattern P of length m within a text T of length $n \geq m$ takes $\mathcal{O}(n - m + 1)$ time on average when P and T are random strings, but this is a completely unreasonable assumption. (Chapter 32 in CLRS4)

Probabilistic analysis helps gain insight into a problem and helps design an efficient algorithm for it, when we have a reasonable assumption on the distribution of the inputs.

A *randomised algorithm* (as opposed to a deterministic algorithm) itself makes random choices, independently of the actual distribution of the inputs. We refer to the time of a randomised algorithm as *expected time* (not: average time).

Examples

- A randomised algorithm by Karger-Klein-Tarjan (1993) computes in $\mathcal{O}(V + E)$ expected time a minimum spanning tree (MST) of a connected undirected graph with vertex set V and edge set E . (Chapter 21)
- A randomised algorithm computes in $\mathcal{O}(m)$ expected time a prime number larger than m for fingerprinting in the Rabin-Karp string matcher. (Chapter 32)

Many randomised algorithms have no worst-case input!

In an *amortised analysis*, we compute the worst-case time of a *chain* of data-structure operations, and average it over the operations. We refer to this time as an *amortised time* (as opposed to an average-case time, as no probability is used here, and to the possibly non-tight worst-case time).

Examples

- A chain of m find-and-compress-paths or union-by-rank operations on disjoint sets of n items takes $\mathcal{O}(m \cdot \lg^* n)$ time, where $\lg^* n \leq 5$ in practice. (Chapter 19)
- In a Fibonacci heap of n items, extracting a minimum takes $\mathcal{O}(\lg n)$ amortised time, and decreasing a key takes $\mathcal{O}(1)$ amortised time. (online chapter; not in AD2)
- Prim's MST algorithm takes at worst $\mathcal{O}(E + V \lg V)$ time when using a Fibonacci heap. (Chapter 21)

Dealing in polynomial time with (instances of) optimisation problems where brute-force or exact solving is too costly:

- A *greedy algorithm* builds a feasible solution decision variable by decision variable, making locally optimal choices in the hope of reaching an optimal solution. Greedy algorithms build either provably optimal solutions (for example, Prim's MST algorithm and Dijkstra's single-source shortest paths algorithm) or at-best optimal solutions.
- A *local search algorithm* repairs a possibly infeasible candidate solution, by reassigning some decision variables at every iteration, until an allocated resource (such as an iteration count or a time budget) is exhausted, in the hope of reaching a feasible or even optimal solution.
- An *approximation algorithm* for an NP-hard optimisation problem builds a feasible solution whose objective value is provably within a known factor of the optimum.

All techniques are orthogonal: there exist randomised local search algorithms, greedy approximation algorithms, etc.

In order to pass, the student must be able to:

- analyse NP-completeness of an algorithmic problem;
- use advanced algorithm analysis methods, such as amortised analysis and probabilistic analysis;
- use advanced algorithm design methods in order to approach hard algorithmic problems in a pragmatic way, such as by using:
 - ▶ randomised algorithms: universal hashing, ...
 - ▶ approximation algorithms
 - ▶ stochastic local search: simulated annealing, tabu search, ...
 - ▶ mixed integer programming (MIP)
 - ▶ Boolean satisfiability (SAT)
 - ▶ SAT modulo theories (SMT)
- present and discuss topics related to the course content orally and in writing with a skill appropriate for the level of education written reports and oral resubmissions!

Course Organisation and *Suggested* Time Budget

Period 3: January to March, budget = 133.3 hours:

- 12 **lectures**, including a **mandatory** guest lecture, budget = 21 hours
- 2 **assignments** with 3 *help sessions*, 1 *grading session*, 1 *solution session* per assignment, on 2 problems each, on *non-exam* topics, to be done by student-chosen duo team: *suggested* budget = average of 30 hours / assignment / student (2 credits)
- 1 written ***closed-book*** **exam** of 3 hours, to be done individually: *suggested* budget = 52 hours (3 credits)
- **Prerequisites:** Algorithms and Data Structures 2 (AD2) (course 1DL231) or equivalent

Examination

Modelling / programming, experimenting, and reporting:

- Mixed integer programming (MIP): Assignment 1
- Stochastic local search (SLS): Assignment 1
- Boolean satisfiability (SAT): Assignment 2
- SAT modulo theories (SMT): Assignment 2

Theory questions, drawn from a published list of potential exam questions:

- Amortised analysis and probabilistic analysis: exam
- Randomised algorithms: exam
- NP-completeness: 50% threshold at exam
- Approximation algorithms: exam

2 Assignment Cycles of 3 Weeks

- *help session a*: attendance strongly recommended!
- *help session b*: attendance strongly recommended!
- *help session c*: attendance strongly recommended!
- Within 10 days of the deadline at 16:00: your *initial score* $a_{ij} \in 0..5$ points for each Problem j of Assignment i , with $j \in 1..2$
- after publication teamwise oral *grading session* on *some* Problems j where $a_{ij} \in \{1, 2\}$: possibility of earning 1 extra point for your *final score*; otherwise final score = initial score
- *solution session* and *help session a*

2 Assignment Credits and Overall Influence

Let a_{ij} be your *final score* on Problem j of Assignment i , with $i, j \in 1..2$:

- 20% threshold: $\forall i, j \in 1..2 : a_{ij} \geq 20\% \cdot 5 = 1$ (< 3) You may not catastrophically fail on individual problems
- 30% threshold: $\forall i : a_i = a_{i1} + a_{i2} \geq 30\% \cdot (5 + 5) = 3$ (< 5) You can partially fail on individual problems or entire assignments
- 50% threshold: $a = a_1 + a_2 \geq 50\% \cdot 2 \cdot (5 + 5) = 10$ The formula for your *assignment grade* in 3..5 is at the course homepage
- Worth going full-blast: Your *assignment score* a is meshed with your *exam score* e in order to determine your *overall course grade* in 3..5, *if* $10 \leq a \leq 20$ **and** $10 \leq e \leq 20$: see the formula at the course homepage

Caution!

- There is a huge jump from AD2 (or equivalent) — with its mostly (pseudo-)polytime algorithms — to AD3, where only NP-hard problems are considered.
- Correctness is required (unlike in AD2), but very easy to achieve with the help of our provided polytime solution checkers or some revealed optima: we grade for speed (and memory usage).
- Especially the MIP, SAT, and SMT modelling tasks are totally unlike anything most of you have ever seen, and this takes time to wrap one's head around.
- Ease or success with the assignments in AD2 does not imply the same ease or the same level of success with the assignments in AD3: the help sessions are strongly recommended, and there is almost no internet help.

Assignment Rules

Register a *team* by Sunday 25 January at 23:59 on Studium:

- Duo team: Two consenting teammates sign up
- Random partner? join the group I want to be randomly assigned.

If you don't register for a team or ask to be randomly assigned by the deadline, then I will assume that you are not participating in the course.

Other considerations

- Teammate scores may differ if no-show or passivity at grading session
- No freeloader: Implicit honour declaration in reports that each partner can individually explain everything; random checks will be made by us
- No plagiarism: Implicit honour declaration in reports; extremely powerful detection tools will be used by us; suspected cases of using or providing must be reported! No use of LLMs or coding assistants. You need to learn the material yourself.

What To Do Now?

- Bookmark and read the entire AD3 website¹, especially the FAQ list.
- Get started on Assignment 1 and have questions ready for the first help session, which is on the 27th January.
- Register a duo team by Sunday 25 January at 23:59.
- Install AMPL (see Studium for a free download with the classroom license) on your own hardware, if you have any.

¹<https://ad3-uu-se.github.io/>