

Industrial Vehicle Routing

Chief Scientist, Dr. scient. Geir Hasle
SINTEF ICT, Oslo, Norway

Invited talk at

Operations Research 2008

Augsburg, Germany, September 3 2008

Outline

- Background and Context
- The Vehicle Routing Problem
- Industrial Aspects
- SPIDER - A Generic VRP Solver
- Ongoing and Further Research

Outline

- Background and Context
- The Vehicle Routing Problem
- Industrial Aspects
- SPIDER - A Generic VRP Solver
- Ongoing and Further Research



Technology for a better society

- An Independent Multi-Disciplinary Contract R&D Organisation
- Established in 1950
- Among the Largest CRO in Europe

Vision:

- Technology for a better Society

Business Concept:

- To meet the needs for Research-Based Innovation and Development for the Private and Public Sectors

1.300 employees in Trondheim



● *Trondheim*

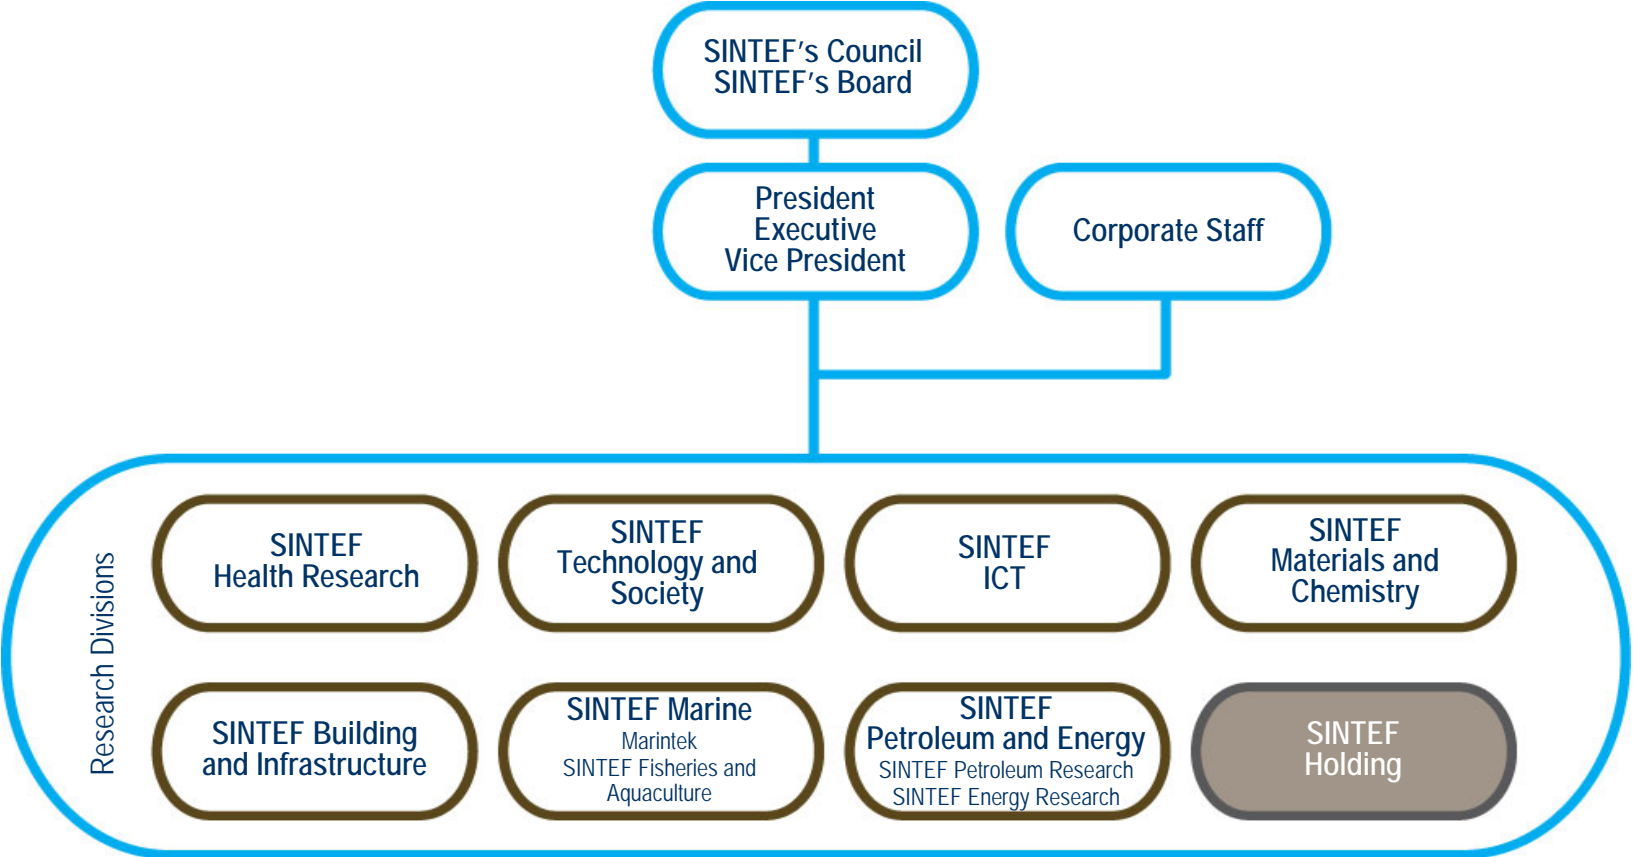
450 employees in Oslo



SINTEF - A Norwegian Contract Research Institute

- Science and engineering – social sciences, health care
- Connected with
 - The Norwegian University of Science and Technology, Trondheim
 - The University of Oslo, Faculty of mathematics and natural sciences
- > 90 % of turnover from industry and public sector contracts
- Annual turnover ~ 250 M€
- Activities
 - Strategic, long term, basic research
 - Contract Research
 - Consultancy
 - Commercialization and spin-offs

The SINTEF Group



Outline

- Background and Context
- **The Vehicle Routing Problem**
- Industrial Aspects
- SPIDER - A Generic VRP Solver
- Applications
- Challenges
- Ongoing and Further Research

The Vehicle Routing Problem (VRP)

Given a fleet of vehicles and a set of transportation orders, find a minimum cost routing plan. That is: allocate each order to a vehicle, and for each vehicle, sequence the stops.

- The VRP central to efficient transportation management
- Applications
 - Distribution or pick-up of goods
 - Dial-a-ride
 - Municipal services
 - Repairman problem
 - Newspaper distribution
 - Waste management
 - Gritting, snow clearing
- Very hard, discrete optimization problem

The Capacitated VRP (CVRP)

- Graph $G=(N,A)$

- $N=\{0,\dots,n\}$ Nodes
- 0 Depot, $i\neq 0$ Customers
- $A=\{(i,j): i,j\in N\}$ Arcs
- $c_{ij} > 0$ Arc cost

- Demand d_i for each Customer i

- either all pickups or all deliveries

- V set of (identical) Vehicles, each with Capacity q

- Goal



- Construct a set of minimum cost Routes that start and finish at the Depot.
- Each Customer shall be visited once (no order splitting)
- Total Demand for all Customers serviced less than Capacity q
- Cost: Total Arc cost (also # Routes in richer VRP variants)

- DVRP – constraint on route length

- VRPTW – VRP with time windows

- VRPB – VRP with backhauling (deliveries first, then pickups)

- PDP – pickup and delivery without going through the depot



Mathematical formulation of VRPTW (vehicle flow formulation)

minimize $\sum_{k \in V} \sum_{(i,j) \in A} c_{ij} x_{ij}^k$ (1) minimize cost

subject to:

$\sum_{k \in V} \sum_{j \in N} x_{ij}^k = 1, \quad \forall i \in C$ (2) each customer once

$\sum_{i \in C} d_i \sum_{j \in N} x_{ij}^k \leq q, \quad \forall k \in V$ (3) capacity

$\sum_{j \in N} x_{0j}^k = 1, \quad \forall k \in V$ (4) k routes out of depot

$\sum_{i \in N} x_{ih}^k - \sum_{j \in N} x_{hj}^k = 0, \quad \forall h \in C, \quad \forall k \in V$ (5) flow balance for each customer

$\sum_{i \in N} x_{i,n+1}^k = 1, \quad \forall k \in V$ (6) k routes into depot (redundant)

$x_{ij}^k (s_i^k + t_{ij} - s_j^k) \leq 0, \quad \forall (i,j) \in A, \quad \forall k \in V$ (7) start of service and driving time

$a_i \leq s_i^k \leq b_i, \quad \forall i \in N, \quad \forall k \in V$ (8) start of service within TW

$x_{ij}^k \in \{0,1\}, \quad \forall (i,j) \in A, \quad \forall k \in V$ (9) arc (i,j) driven by k

The VRP is very hard, computationally

- Belongs to class of strongly **NP**-hard optimization problems
- Computing time for **all VRP optimization algorithms** grows "exponentially" with the number of nodes (probably ...)
- Exact methods are "easy" to design, may take forever
- Currently, exact methods for the C/DVRP can **consistently** solve instances with **some 70 customers only** to optimality within a few minutes
- For generic industrial VRP tools, we must give up the quest for optimality and resort to some form of approximation
- What is the optimal plan, anyway?

VRP in Operations Research

- since 1959
- thousands of papers
- mostly generic work on idealized VRP models
- still, one of the successes of OR
- a tool industry has emerged, based on research results
- tremendous increase in ability to "solve" VRP variants
 - general increase of available computing power
 - methodological improvement
- still much to do, VRP research has never been more active
- short road from research results to industrial benefits

VRP in Operations Research

- General, idealized formulations
- Extensions studied in isolation
 - time windows
 - multiple depots
 - inventory constraints
 -
- Very fruitful for
 - understanding VRP variants
 - developing highly targeted algorithms
- Not always relevant to real-life problems
- Some application specific work
- Recently: More holistic approach: "Rich VRPs"
 - General model, many aspects of industrial problems
 - Robust algorithms

Extensions in the VRP-literature

- Location Routing
- Fleet Size and Mix
- VRP With Time Windows
- General Pickup and Delivery
- Dial-A-Ride
- Periodic VRP
- Inventory Routing
- Dynamic VRP
- Capacitated Arc Routing Problem

- All kinds of algorithmic approaches
 - exact
 - approximation methods
 - based on systematic, exact methods
 - metaheuristics



LRP
FSMVRP
VRPTW
GPDP
DARP
PVRP
IRP
DVRP
CARP

Outline

- Background and Context
- The Vehicle Routing Problem
- **Industrial Aspects**
- SPIDER - A Generic VRP Solver
- Ongoing and Further Research

Industrial aspects – VRP tool

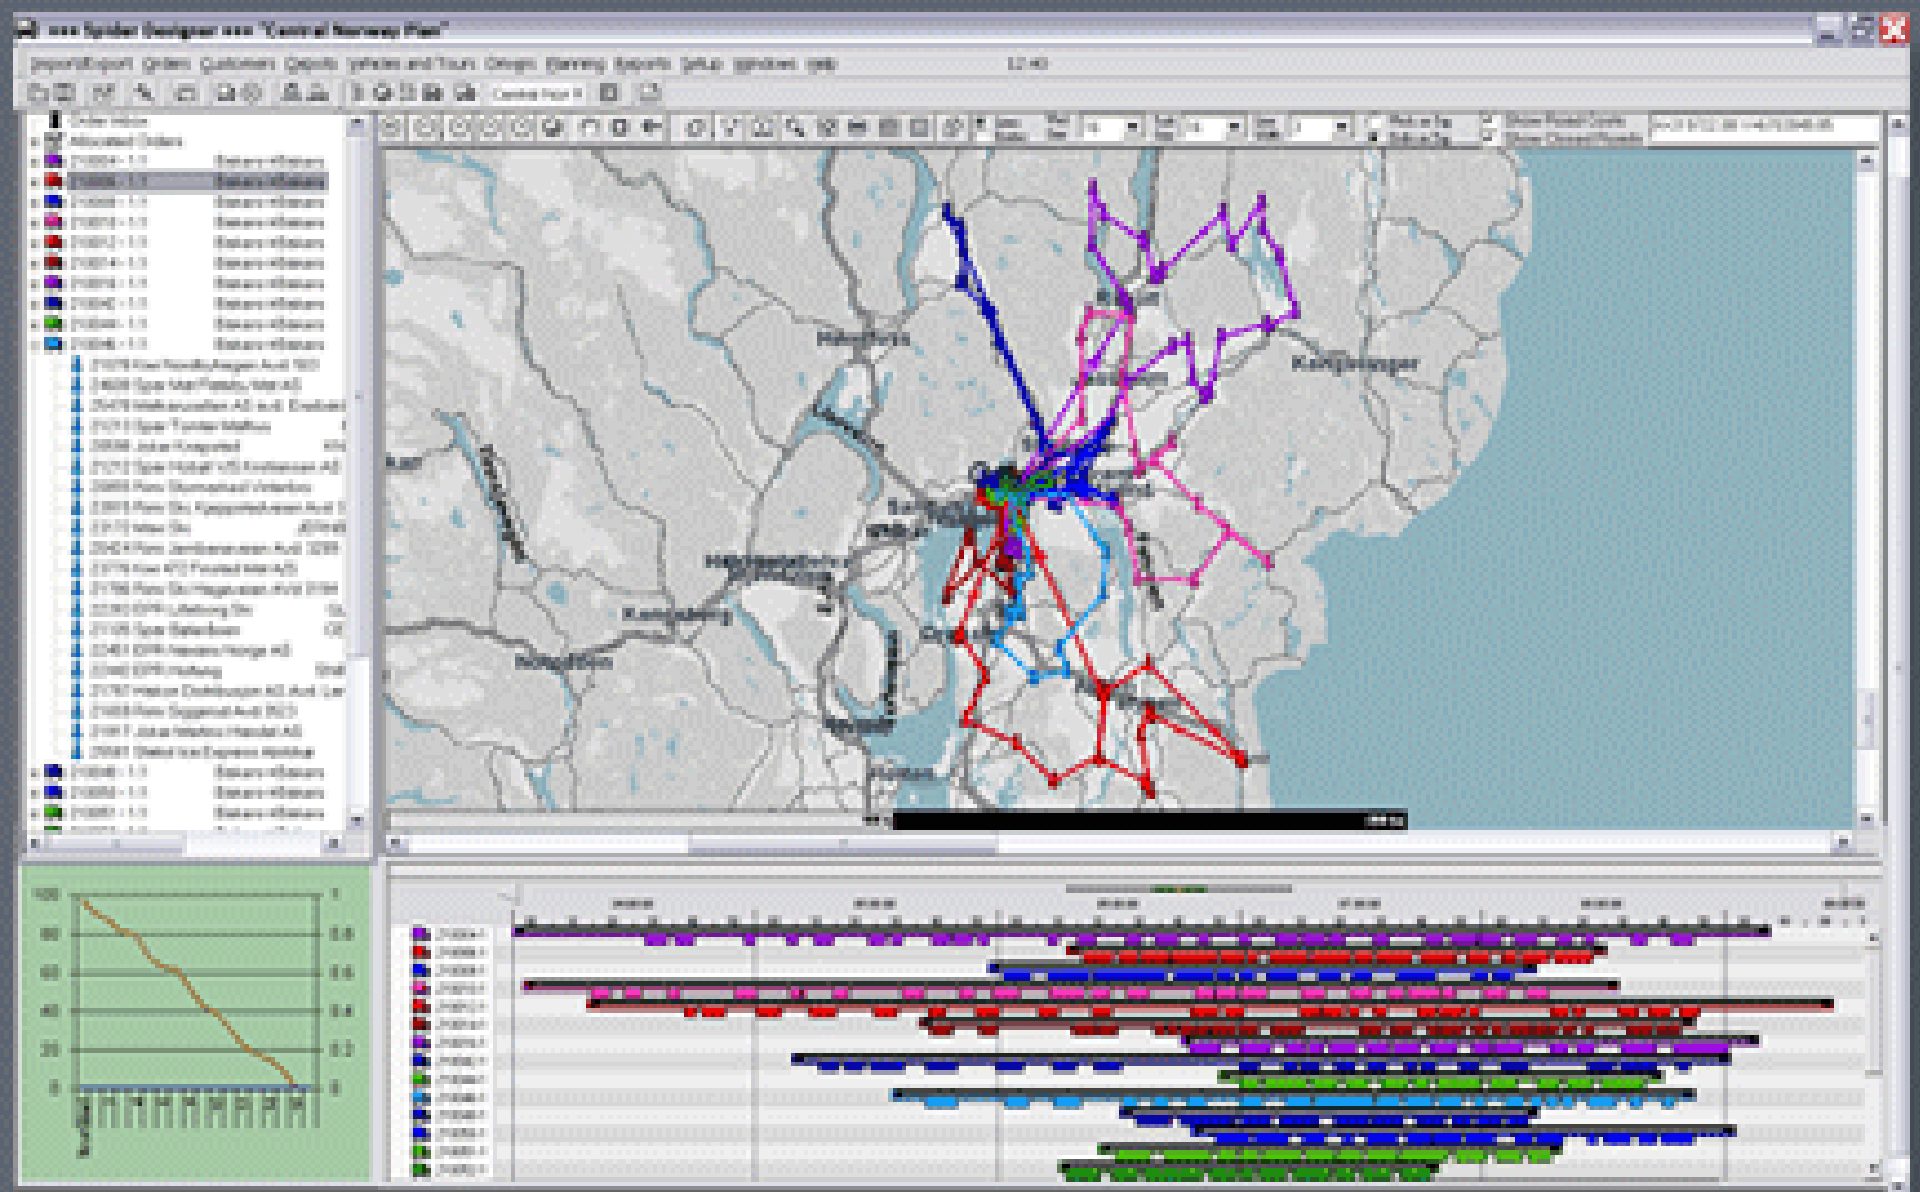
- Adequate model of the application
- High quality information
 - addresses
 - distances, times, costs
 - orders
 - fleet and driver information
- Manual planning
- Ease of use
- Integration with ERP etc.
- High quality solution in short time – powerful VRP solver

Industry needs rich VRP models

- Type of service and operation
 - multiple depots, no depot
 - different order types (delivery, pickup, direct, service, ...)
 - order splitting, flexible order volumes
 - arc routing
 - multiple tours per day
 - periodic problems
 - connection to inventory and manufacturing
- Constraints
 - capacity, several dimensions, hard and soft
 - time windows, multiple, hard and soft
 - precedences
 - (in)compatibilities
- Realistic distances, times, costs
- Cost components
 - multiple criteria
 - soft constraints and penalties
- Uncertainty and Dynamics

Research on Rich VRPs & related problems at SINTEF

- Industrial Contracts since 1995
- Strategic Research
 - European Commission FP III, IV, V (e.g., GreenTrip 1996-1999)
 - Norwegian Research Council (RCN)
 - Internal Projects, students
- Generic VRP Solver - SPIDER (1995→)
 - Commercialization from 1999
 - GreenTrip AS → SPIDER Solutions AS
- TOP Programme 2001-2004 (RCN) <http://www.top.sintef.no/>
 - Basic Research on Rich VRP and related problems
 - VRPTW
 - Shortest Path Problem in Dynamic Road Topologies
- DOMinant 2005-2009 (RCN)
 - Basic Research on Rich VRP, together with NTNU and Molde University College
 - Inventory routing and fleet composition, road-based and maritime applications
- Innovation Projects supported by Research Council of Norway
 - “I Rute” (2001 – 2004) Bulk transportation
 - “DOiT” (2004 – 2007) Stochastic and Dynamic Routing
 - “EDGE” – (2005 – 2008) Huge Scale Routing
 - “Effekt” – (2008 – 2011) Huge Scale Routing
- Ship routing – many ongoing projects




SPIDER Designer - Applications

- Distribution of bread (Bakers)
- Mail collection and distribution (Posten Norge)
- Local pickup and delivery (Schenker)
- Newspaper distribution, 1st tier (Aftenposten, Dagbladet)
- Newspaper distribution, last mile (Aftenposten, Stavanger Aftenblad)
- Collection of milk from farms (TINE)
- Distribution of fodder to farms (Landbruksdistribusjon)
- Distribution of fuel oil (Hydro Texaco)
- Location analyses, depot (obnoxious facility location, Norsk Gjenvinning AS)
- Distribution of blood (Ullevål sykehus)
- Distribution of groceries (REMA 1000)
- Distribution of magazines (Bladcentralen)
- Distribution of ice cream (Diplom Is; Hennig Olsen)
- Dial-a-ride, elderly, hospital patients (Nor-Link)

- Savings 5-35%, depending on application

Products - architecture


SPIDER Solutions AS



Distribution Innovation AS



Geomatikk AS



Outline

- Background and Context
- The Vehicle Routing Problem
- Industrial Aspects
- **SPIDER - A Generic VRP Solver**
- Ongoing and Further Research

SPIDER - A Generic VRP Solver

- Designed to be widely applicable
- Based on generic, rich model
- Predictive route planning
- Plan repair, reactive planning
- Dynamic planning with stochastic model

- Framework for VRP research


SPIDER - Generalisations of CVRP

- Heterogeneous fleet
 - Capacities
 - Equipment
 - Arbitrary tour start/end locations
 - Time windows
 - Cost structure
- Linked tours with precedences
- Mixture of order types
- Multiple time windows, soft time windows
- Capacity in multiple dimensions, soft capacity
- Alternative locations, tours and orders
- Arc locations, for arc routing and aggregation of node orders
- Alternative time periods
- Non-Euclidean, asymmetric, dynamic travel times
- Compatibility constraints
- A variety of constraint types and cost components
 - driving time restrictions
 - visual beauty of routing plan, non-overlapping

Order

■ Different types:

- Delivery
- Pickup
- Direct (P&D)
- Service




■ Plan structure

- task
- task sequence
- tour
- plan



Tour

- Time windows and locations given by start/stop tasks
- Selected vehicle and driver (alternative equipages)
- Linked tours
 - (may have different vehicles)




Vehicle and Driver

- Capacity
- Travelling attributes
 - Speed profile
 - Height, weight, length
 - Obey one-way restrictions?
- Driving time regulations

Cost elements

- Travel cost (distance, time, tolls)
- Tour usage cost
 - Cost for starting a tour
 - Cost per order on tour
- Cost for unserviced order
- Waiting time cost
- Cost for alternative locations
- Cost for same/different location
- Cost for breaking work regulations
- Cost for “ugly”, overlapping routes

Constraints



■ Consistency

- Complete order
- Pickup/delivery (Direct orders) same tour and precedence

■ Time

- Travel time, Duration
- Time windows, multiple, hard and soft

■ Vehicle capacities, multiple dimensions, hard and soft

■ Total capacities over a set of tours

Constraints

- Compatibility vehicle/order vehicle/location product/compartment
- Orders on same tour
- Corresponding locations (when alternatives)
 - Order: Choose corresponding locations for pickup and delivery task
 - Tour: Choose corresponding locations for start/stop tasks
- Corresponding time periods for sets of orders
 - E.g. Delivery day 1,3,5 or day 2,4,6

Locks

- Prevent optimiser changing part of plan
- Task: Time lock
- Tour: Lock whole or initial part of tour
- Order: Lock (un)assignment

Uniform Algorithmic Approach

■ Goals

- Reach a good local optimum fast
- Explore interesting parts of search space efficiently

■ 3 phases

- Construction
- Tour Reduction
- Iterative Improvement: Iterated Variable Neighborhood Descent

■ based on

- Variable Neighborhood Descent (Hansen & Mladenovic)
- Diversification when VND reaches local optimum
- Iterated Local Search (Martin, Lourenço et al)

Construction of Initial Solution

- Various Sequential Construction Heuristics
- Extended to cover Richer Model
 - Several types of order
 - Non-standard constraints
 - Non-homogeneous fleet
 - Multiple Depots
 - Multiple Tours per Vehicle
- New Constructor
 - Inhomogeneous problems
 - Multiple depots
 - Multiple tours per vehicle
 - More search


Variable Neighborhood Descent

- Repertoire of 12 operators

- Insert
- Remove
- 2-opt
- Or-opt
- 3-opt
- Relocate
- Cross (2-opt*)
- Exchange (Cross-Exchange)
- Tour Depletion
- Change alternative location
- Change alternative time period
- Change vehicle
- Variety, efficiency, sequence?

Illustration: Exchange

- Inter-tour operator
- Full neighborhood is typically large
- Remedy: Limit on maximum segment length
- Alternative: Focus on promising cut points



Iterated Local Search (Martin, Lourenço et al)

- Goal: Efficient search in new basins of attraction
- When VND reaches local optimum: diversify
- Random restart
- Alternative initial solution
- Path relinking
- Noising
- LNS, VLNS
- Change of objective

Large Neighborhood Search (Shaw)

- Take away a substantial number of commitments
 - randomly
 - “similar” commitments
- Reconstruct with fast insertion method
 - Cheapest insertion
 - Regret-based insertion
- Accept new solution if
 - better, diverse
 - Threshold Acceptance
 - Simulated Annealing
- Iterate
- Alternative modifications
 - Limited Local Search
 - Full VND Machinery
 - Distance Metrics

Computational Experiments

- Standard test problems in literature for most VRP variants
- "World championship" in Vehicle Routing
- For most instances, the optimal solution is not known
- Testing against the best methods reported in the literature
- "Unfair" comparison: competing with solvers optimized for solving a specific, idealized problem
- Nice to know how well you perform ...

Experimental studies PDPTW

<http://www.top.sintef.no/>

- 354 standard test problems Li & Lim: 100 – 1.000 orders

Author	100	200	400	600	800	1000	Total
Li & Lim	41	14	6	2	0	0	63
SINTEF	13	11	5	4	9	4	46
BVH	2	8	6	3	4	1	24
TS	0	1	2	0	0	2	5
SR	0	26	41	51	47	51	216
Total	56	60	60	60	60	58	354

- BVH: Bent & Van Hentenryck: Two-stage, Hybrid Local Search
 - 1. Minimize # tours by SA, modified objective
 - 2. Minimize Distance with Large Neighborhood Search
- TS: Tetrasoft, Danish tool vendor, Method unknown
- SR: Stefan Röpke, Univ. Copenhagen
 - Uniform approach
 - Adaptive Large Neighborhood Search

Results on CVRP

- Rounded distances (1)

■ Augerat et al

- 74 instances, 15-100 customers, limit on # vehicles
- A (27): 0.15% from optimum, 19 optimal, 0-541 s, average 100 s
- B (23): 0.11% from optimum, 17 optimal, 1-582 s, average 95 s
- P (24): 0.23% from optimum, 15 optimal, 0-450 s, average 99 s

■ Christofides & Eilon (1969)

- 13 instances, 12-100 customers, limit on # vehicles
- E (13): 0.32% from optimum, 2 optimal, 7 - 487 s, average 178 s

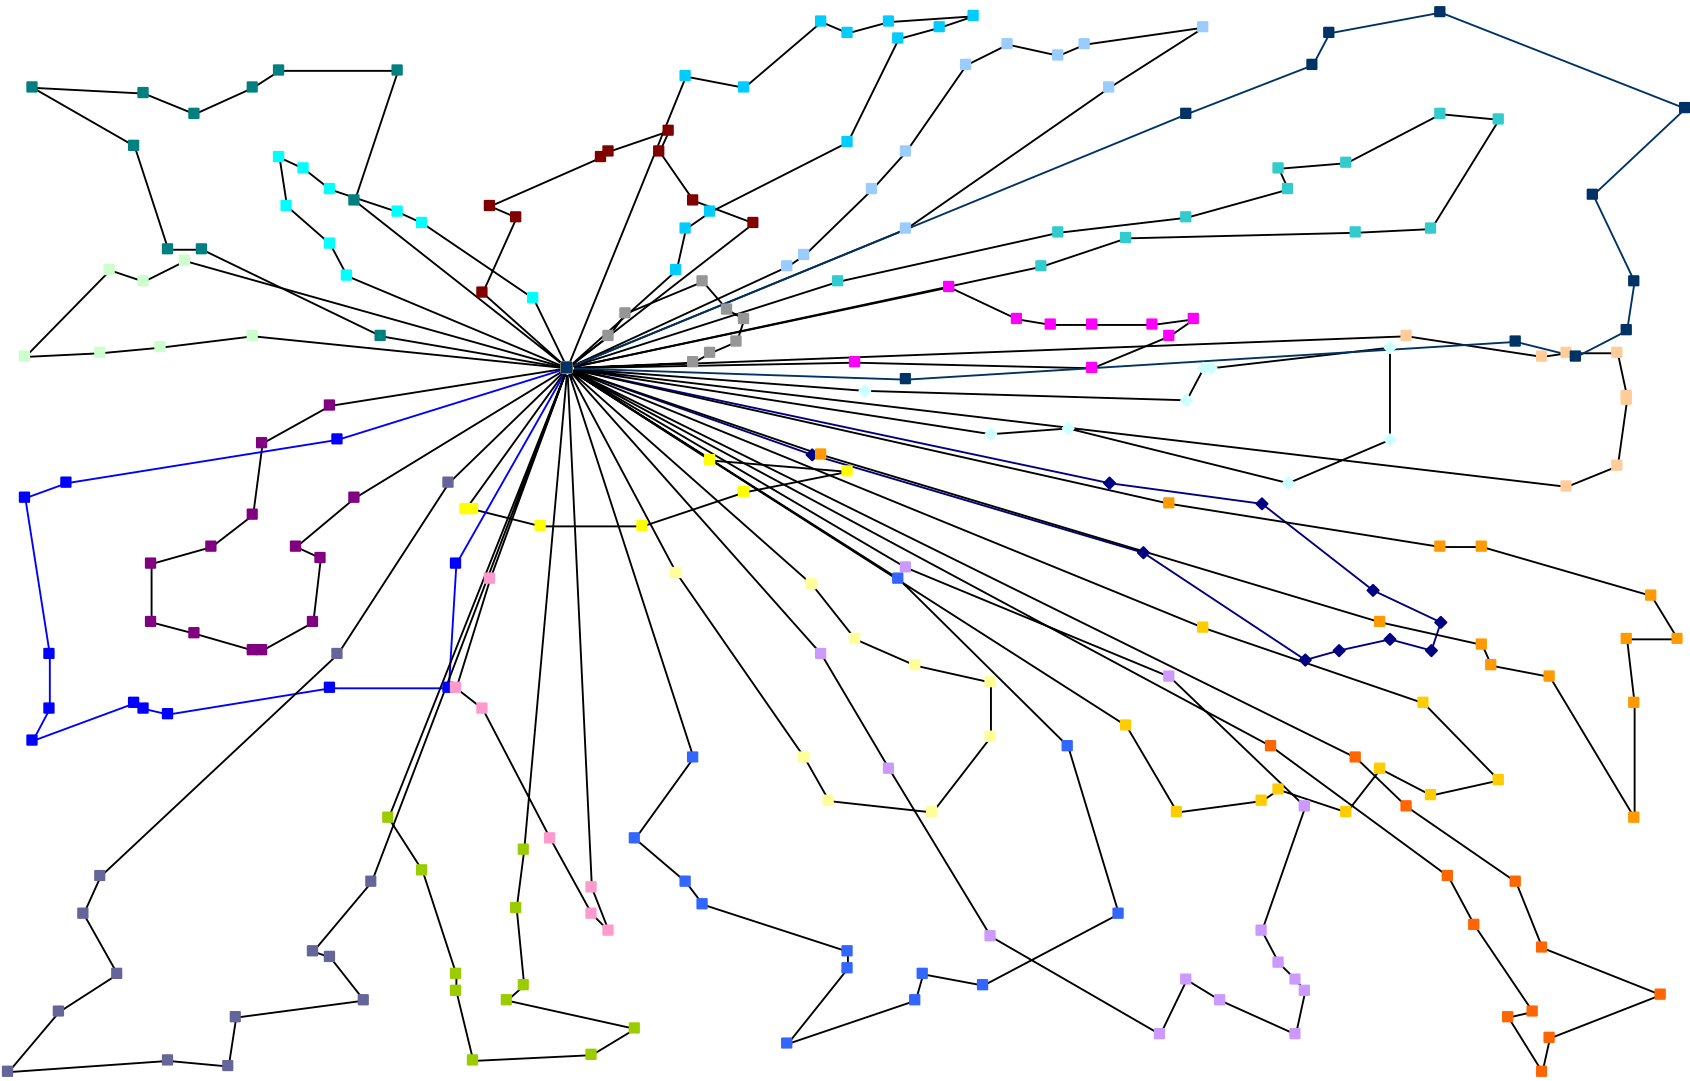
■ Fisher

- 3 instances, 44 - 134 customers, limit on # vehicles
- F (3): 0.27% from optimum, 7 optimal, 0-289 s, average 103 s

■ Gillett & Johnson, 1 instance: **G-n262-k25**

- best known, -7.09% from previous best known (5685)

G-n262-k25: 5685 vs. 6119, 5767 CPU s




Results on CVRP

- Rounded distances (2)

- Christofides, Mingozzi & Toth (1979)
- 5 CVRP instances, 100-199 customers, limit on # vehicles
- 0.99 % from optimum, 21-1255 s, average 404s
- 1 optimal
- 1 new best known -1.6 %
- 1 worse by 2.7 %
- 1 worse by 2.9 %
- M-n200-k16: no feasible solution previously known
- Feasible solution found
 - after 533s: distance 1371

M-n200-k16: First known feasible solution



Outline

- Background and Context
- The Vehicle Routing Problem
- Industrial Aspects
- SPIDER - A Generic VRP Solver
- Ongoing and Further Research

Challenges

- Instance robustness
 - good performance on all kinds of instances
- Scalability
 - very large size problems
- Extendability
 - new operators, new VND sequence

Scalability – Solving "Huge" VRPs

- Faster VRP solver
 - heuristic investigation of neighborhoods
 - clever sequencing of constraint and objective component evaluation
 - more powerful metaheuristics
 - hybrid methods, even combining exact methods and heuristics
 - exploiting parallel computers, multiple cores
- Problem Decomposition
 - geographically
 - temporally
 - based on product
 - clustering methods will be useful
- Abstraction and Aggregation
 - clustering methods will be useful
- Multi-level search and collaborative solvers

Fleet composition and minimization

2 phase tour reduction approach

- Morten Smedsrud, ongoing work
- Alternates between 2 phases, performing a local search with fast neighborhood operators in each phase.
- Operators are: 2-opt, 3-opt, or-exchange, exchange, cross.
- 1st phase uses a pure tour reduction objective
- 2nd phase uses a traditional minimize number of vehicles first, tour distance second.
- The phases are otherwise identical

Tour reduction objective

$$o(\sigma) = \left\langle |\sigma|, -\sum |r|^2 + N * st(\sigma), mdl(\sigma) \right\rangle$$

- Based on the tour reduction phase in Bent and Van Hentenryck's (BVH) VRPTW paper, Transportation Science Vol. 38, No. 4, November 2004, pp. 515–530
- Primary objective is number of tours
- Secondary objective is maximizing the sum of the square “length” of the tours as in BVH. This is modified relatively to BVH by subtracting the length of the shortest tour multiplied by number of orders in the case from the square sum.
- Motivation of the modification of secondary objective is so that it is always preferred to make the shortest tour shorter rather than moving an order from a long tour to even longer tour.
- Tertiary objective is minimizing the minimum delay in the plan, First introduced in Homberger and Gehring's VRPTW INFOR paper of 1999.

Observations

- When starting a new local search we observe that the operators initially find a lot of improvements to the plan but slows down eventually until a new local minimum is reached.
- After switching objectives operators find improvements fast again.
- Running to local minimum then switching between objectives seems to yield good results in tour reduction.
- Seems to give good diversification, quality of initial solution seems irrelevant.
- Often the local minima after running with tour reduction phase ended with shortest tours consisting of only 1-3 orders, so a tour depletion procedure was implemented

Tour depletion

- Try to shorten each of the shortest tours by moving orders to the remaining tours.
- For each of the remaining tours try to move orders to the other remaining tours to make room for orders from the shortest tour.
- Can be applied recursively, but computation time increases geometrically so we only do one step for now.
- If a move from the shortest tour is possible it is performed, and we continue with the new solution.
- If no move is possible, backtracks and tries move to next remaining tour.


Results on H&G VRPTW 400

- Total number vehicles 1391, a lot closer to the best known to us result of 1388 (Prescott-Gagnon et al. 2008).
- Total distance 409868. Compared to 390771 in P-G.
- Maximum amount of computational time in tour reduction phase before finding plan with fewest vehicles: 2h22min
- Average time before finding the plan with fewest vehicles: 17.5min

Summary

- The Vehicle Routing Problem is a key to efficient transportation
- The VRP is a very hard optimization problem
- Thousands of VRP papers, mostly generic, idealized
- OR success story, tool industry
- Industry needs rich models, powerful algorithms
- Still challenges, VRP research has never been more active
- Generic VRP Solver SPIDER
 - rich, generic model
 - resolution algorithm based on metaheuristics
 - results
 - ongoing work
- Future research topics

**Geir Hasle and Oddvar Kloster: *Industrial Vehicle Routing*.
Chapter (pp 397-435) in
Hasle, Geir; Lie, Knut-Andreas; Quak, Ewald (Eds.)
Geometric Modelling, Numerical Simulation, and Optimization:
Applied Mathematics at SINTEF
2007, XI, 558 p. 162 illus., 59 in color., Hardcover. ISBN: 978-3-540-68782-5
<http://www.springer.com/>**



Herzlichen Dank für Ihre Aufmerksamkeit!

Industrial Vehicle Routing

Chief Scientist, Dr. scient. Geir Hasle
SINTEF ICT, Oslo, Norway

Invited talk at

Operations Research 2008

Augsburg, Germany, September 3 2008