VEHICLE ROUTING WITH SPACE- AND TIME-CORRELATED STOCHASTIC SPEEDS

STEIN W. WALLACE, NHH NORWEGIAN SCHOOL OF ECONOMICS
ZHAOXIA (LEVIN) GUO, SICHUAN UNIVERSITY
MICHAL KAUT, SINTEF, NORWAY

Vehicle routing

- Basic logistics problem. Thousands of papers and multiple applications.
- A set of customers with demands
- A set of vehicles with capacities
- Travel times between customers
- Time windows (when you can visit customers)
- Service times (how much time will the delivery take?)
- … and many more variants

Important in city logistics

- Major tool for planning the distribution of goods from depots to "customers"
- Could be from major warehouses to shops
- Could be from depots or warehouses to customers, including "last mile".

- Which truck / bicycle takes care of which customers, and what routes should the trucks / bicycles follow?

Uncertainty / stochastics

- Obviously (agree?), when you plan distribution like this, many things are not known to you, depending on ***when*** you make the plan, as well as the specific settings
- Speeds / travel times (traffic conditions)
- Demand (especially when you are collecting or the plan is made beforehand)
- Service times
- Sometimes you do not even know the customers (since they appear while you are driving)

Stochastic vehicle routing

- Some literature over the years
- Mostly on stochastic demand, then travel time (speed), service time and random customers.
- Reviews:
  - V. Pillac, M. Gendreau, C. Gueret and A.L. Medaglia, A review of dynamic vehicle routing problems. *EJOR* 225(1) 2013, 1-11.
  - J. Oyola, H. Arntzen and D.L Woodruff, The stochastic vehicle routing problem: a literature review, Parts I and II –*EURO J Transp and Logist*. (2016 and 2017)
  - U. Ritzinger, J. Puchinger and R.F. Hartl, A survey on dynamic and stochastic vehicle routing problems. *Inter J Production Research* 54(1) 2016, 215-231.

NHH

Random variables

- So people have started to look at randomness, but almost exclusively independent (or uncorrelated) randomness.
- The reason is not that we believe that the phenomena are independent, but that we do not know how to handle dependence in large-scale optimization models.
- The demand for beer depends on the weather across all our customers
- Service times may also depend systematically on the weather
- And travel times on streets in a city are random but obviously positively correlated to different degrees.

Correlated stochastic speeds

- We cover the case where speeds are correlated in time and space:
  - Speeds on the same link in neighboring time periods are correlated; If the speed is low in one period it is probably low also in the next period for a given link.
  - Speeds on links that meet in an intersection are correlated – if there is trouble on one link into an intersection, there is probably also trouble on other links into the same intersection.
- This is new. Old papers that mention “time dependence” simply mean that speeds depend on the time of day. No correlations.

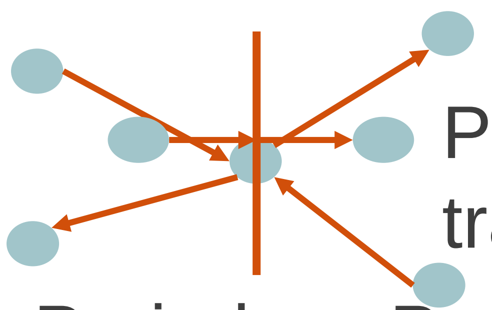

Positively correlated travel times (speeds)

All these links have positively correlated speeds (travel times)

- A map of Beijing with 142 nodes and 418 links (roads).
  - 3 time periods: 418*3 = 1254 correlated random variables
  - 60 time periods: 418*60 = 25,080 correlated random variables

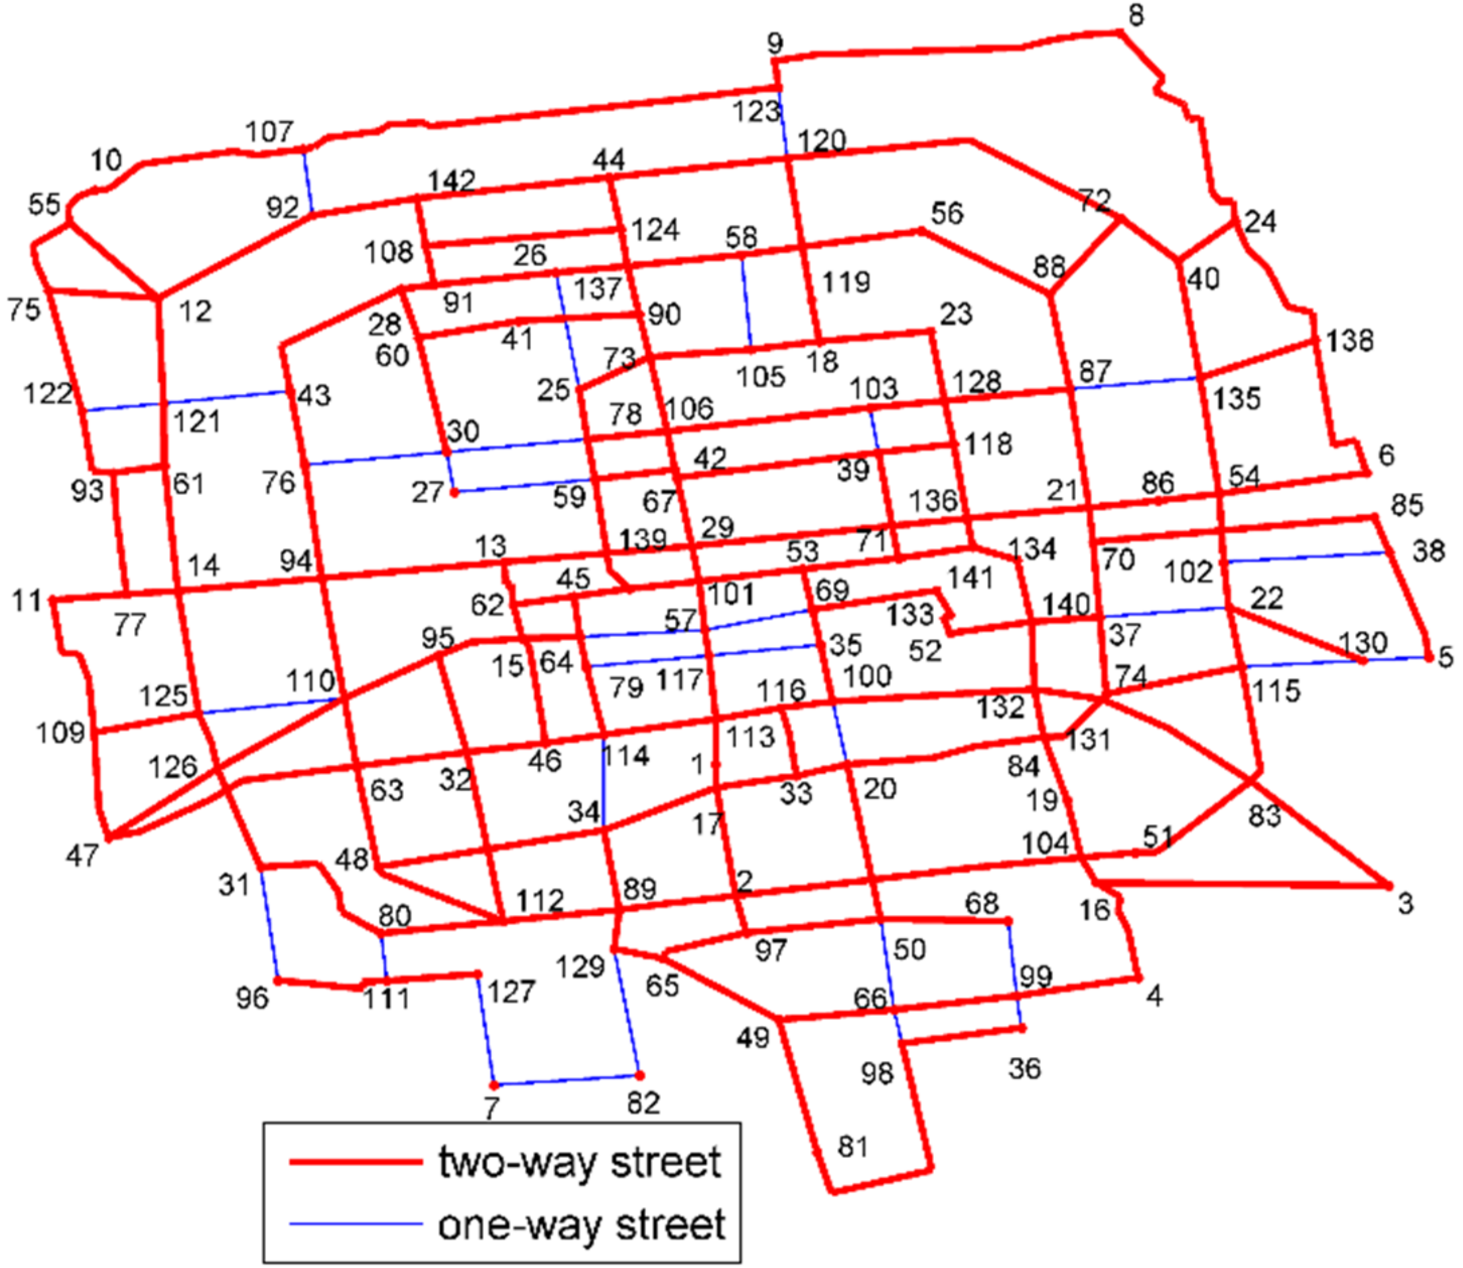

Two-stage model - first example

- Stage 1: Pick routes (sequence of customers) for K capacitated vehicles visiting N customers (living in nodes on the Beijing map) with given demands
- Second-stage: Pick routes between the customers based on realized speeds.
- Goal: Minimize expected excess travel time (over-time pay for drivers).
- The Value of the Stochastic Solution (VSS) is about 10% in this setup.

Modeling correlated speeds

- Normally we use customer-to-customer speeds (travel times) in deterministic models.
- But with stochastic speeds, that does not work.
  - Measures are most likely on road links – especially if they are collected by the system; toll roads, traffic management systems …
  - We might have collected our own data for customer-to-customer speeds, but dependence will be “impossible” to work out as pairs share links – but not always the same ones
  - We would need an enormous number of data points – and for many customer-pair / time-period combinations we would have no data points at all.
  - ***If we assume that customer-to-customer speeds are uncorrelated, we are assuming that no customer pairs share any links.***

- We need to operate on the map

- Huang et al (2017), Time dependent vehicle routing problems with path flexibility. *Transportation Research B* 95:169-195.

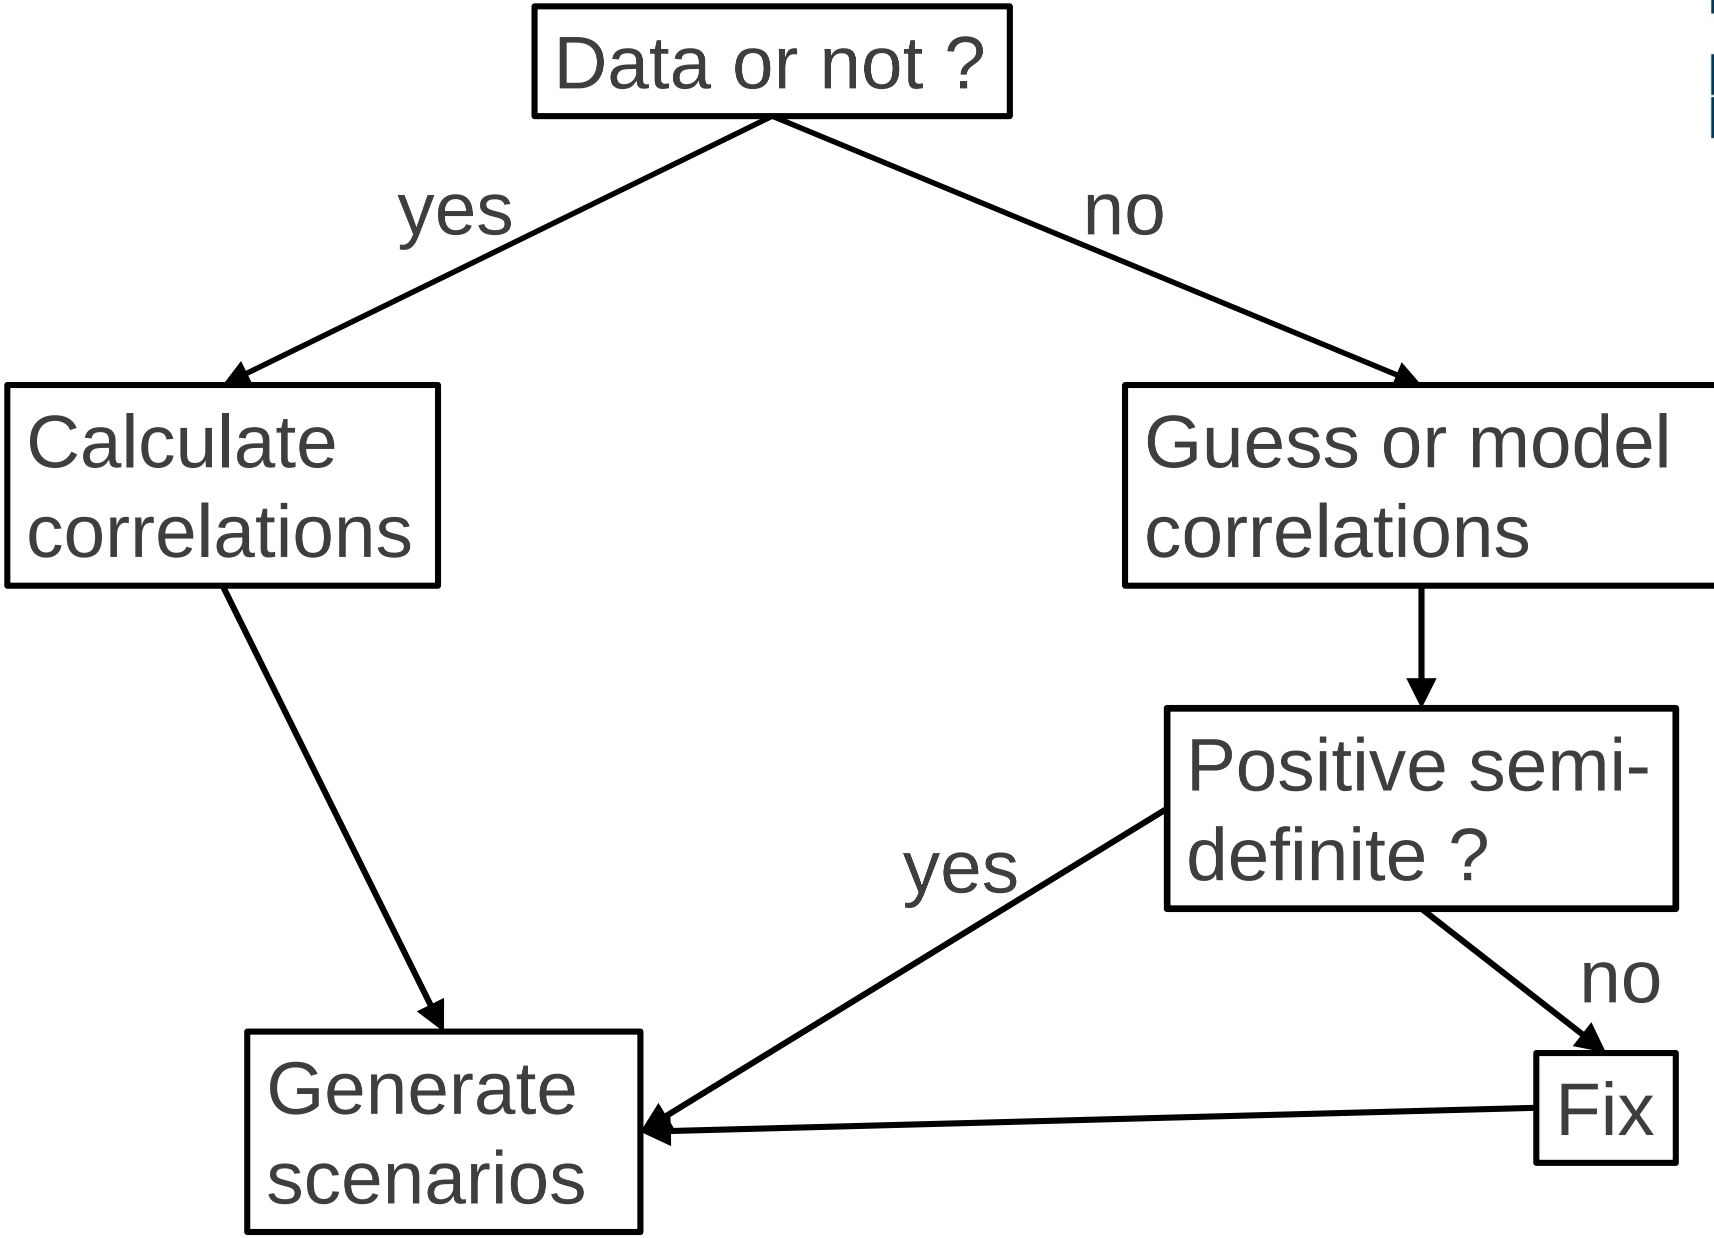
Data or not ?
yes
no
Calculate correlations
Guess or model correlations
Positive semi-definite ?
yes
no
Generate scenarios
Fix

So without data, we can just assume a distribution – right?

- All we need is a 25,000-dimensional dependent random variable
- At least for testing (and academic) settings, just invent one.
- But should we require that what we use makes sense – that the correlations make sense?

- Claim: It is almost impossible (and truly impossible as a general approach) to guess a meaningful distribution of this type.

Correlation matrix

- Defining correlations is less than defining a full distribution
- For example, what about setting up a matrix with
  - The correlation between speeds on the same link in neighboring time periods is α and between neighboring links in the same period β. All others zero.
  - Most likely (but not for sure) this will not lead to a correlation matrix.
- There are ways to «fix» such a matrix to make it positive semi-definite, but that will lead to peculiar correlations, also negative ones, here and there.

With data

- I have not investigated this, so off you go …
- In this kind of dimensions, though a correlation matrix can be calculated, how many measures do we need for ending up with a meaningful one?
- That is, how many measures do we need for the empirical distribution to be useful or meaningful?

Some additional issues of size

- If we have 25,000 dependent random variables, we have over **310 million** distinct correlations to handle (and store …). ***And this problem isn't even particularly big.***
- So even if we had proper data, and could calculate correlations directly, we would be in trouble …
- And fixing a matrix that is not positive semi-definite to make it a correlation matrix is numerically challenging in such dimensions
- Putting the small ones to zero will most likely lead to a matrix that is not positive semi-definite.

- We need scenarios, but what can we do in such dimensions? Gendreau, Jabali and Rei (2016) said that the necessary number of scenarios will "likely be quite large". And they were not thinking of such dimensions.

Scenario generation somewhat generally

- Sampling is always possible if you have something to sample from – which is far from obvious
- It is actually hard to sample from a 25,000 dimensional (dependent) distribution even if you have one
- You could sample from the empirical distribution (if you have one), but will it make sense?

- And there is the issue of how many samples you need in these dimensions. I do not know …

Other than sampling

- I would claim that no available methods would be able to handle, numerically, this kind of dimensions.
- The calculations would kill us … feel free to disagree.

- Remember that we are talking about dependent (correlated) cases.

Our approach to scenarios

- If we cannot define a distribution, not even a meaningful correlation matrix, and in any case could not properly handle it, what can we do?
- We use a scenario generation method from Kaut (2014), based on copulas, that takes as input
  - Marginal distributions
  - Any subset of correlations, and we use
    - The same link in neighboring periods
    - Neighboring links in same period
- This is very well suited for VRPs
  - A route uses exactly such pairs
  - The expected length of any route will be correct
- In the largest example, we specify about 264,000 correlations.

Michal Kaut (2014) A copula-based heuristic for scenario generation, *Computational Management Science* 11 (4) 503-516.

… but

- The method does not control the other correlations. They will often be OK (close to zero) but not always.
- What problems can that cause?
  - The distribution of route lengths will be off if two links with a spurious correlation sit on the same route.
  - And that could affect the optimal routes as methods will be “smart”. (We believe using heuristics actually helps here).

Solution approach

- We assume that the problem is to be solved by a heuristic
  - A part that finds the sequences of customers
  - A part that evaluates any solution
- We shall mostly be concerned about the evaluation part.
- So, in principle, any deterministic heuristic can be applied to the stochastic case.

Evaluating on one scenario

- Starting point: Sequence of customer nodes
- If two customers are neighbors, travel directly
- Otherwise, use present travel speeds to find the shortest path to the next customer (Dijkstra).
- On this path, calculate properly the travel time using speeds in time and space on the scenario.

The number of scenarios

- Even if we have a distribution, there is a question:
  - How many scenarios do we need? (… likely many …)
  - Too few: Just noise
  - Too many: Can’t handle and solve the problem
- And we don’t even have a proper distribution …

- We develop a (new) way to check how many we need – also considering the size and the hardness of the optimization problem

Stability test

- Take a few instances of the problem (in our tests always on the Beijing map).
- For each instance, generate a set of feasible solutions (one coming from the deterministic problem)
- For different *S* (#scenarios), generate 9 scenario trees of size *S-4* up to *S+4* – Kaut (2014) is “almost” deterministic.
- Find the objective function value for each solution and each scenario tree.
- Find the max relative difference for each instance and *S*
- Pick the smallest *S* with acceptable relative difference.

Results of stability tests for the test case with 18 customer nodes, 3 vehicles and 3 time periods.

Number of scenarios		10	30	50	100	500	1000
Solution 1	AOV	0.8640	0.8615	0.8619	0.8616	0.8648	0.8646
	RD	4.99%	2.85%	3.79%	1.65%	0.65%	0.33%
Solution 2	AOV	1.3138	1.3080	1.3080	1.3107	1.3106	1.3105
	RD	4.56%	1.92%	0.91%	0.87%	0.36%	0.14%
Solution 3	AOV	1.7906	1.7811	1.7882	1.7863	1.7864	1.7858
	RD	3.63%	1.33%	1.49%	0.53%	0.32%	0.19%
Solution 4	AOV	4.0130	4.0007	4.0023	4.0068	4.0078	4.0073
	RD	1.58%	1.17%	0.96%	0.57%	0.27%	0.23%
Solution 5	AOV	4.7591	4.7572	4.7522	4.7538	4.7552	4.7563
	RD	2.73%	2.59%	1.51%	0.73%	0.58%	0.30%
Solution 6	AOV	4.7591	4.7572	4.7520	4.7535	4.7548	4.7559
	RD	2.73%	2.59%	1.51%	0.73%	0.59%	0.30%
Solution 7	AOV	3.2539	3.2447	3.2362	3.2350	3.2375	3.2402
	RD	1.62%	1.79%	1.22%	1.19%	0.37%	0.33%
Solution 8	AOV	3.0840	3.0612	3.0672	3.0632	3.0625	3.0624
	RD	1.53%	0.87%	1.35%	1.17%	0.28%	0.27%
Solution 9	AOV	2.8224	2.7840	2.7793	2.7780	2.7845	2.7863
	RD	7.66%	1.85%	2.57%	1.32%	0.55%	0.44%
Solution 10	AOV	4.2755	4.2517	4.2485	4.2382	4.2469	4.2483
	RD	2.80%	1.85%	1.44%	1.04%	0.35%	0.33%

Number of scenarios used and corresponding relative differences

	$P=3$		$P=5$		$P=15$		$P=30$		$P=60$	
(N,K)	S	RD	S	RD	S	RD	S	RD	S	RD
$(18,3)$	100	1.7%	60	0.7%	65	1.0%	45	0.9%	35	1.0%
$(32,5)$	55	0.9%	20	0.9%	25	0.8%	22	0.9%	32	0.8%
$(48,7)$	25	0.8%	15	1.0%	15	0.6%	15	0.9%	15	0.9%
$(64,9)$	15	1.0%	15	0.6%	15	0.6%	15	0.8%	15	1.0%

What is going on?

- In dimension 25,080 we need only 15 scenarios to obtain a relative difference of less than 1%.
- The larger the problem instances (the map is always the same) the fewer scenarios. True both with respect to time periods and number of routes.
- ***Errors caused by strange correlations are part of these error bounds***.
- We need fewer scenarios for larger cases because the errors from incorrect correlations matter less – they are averaged out.
- The errors do not go to zero as *S* increases due to the correlations – but they are small enough.

Stability on *m* – the number of trees is *2m+1*

(N,K,P)	S	$m=4$	$m=6$	$m=8$	$m=10$
(18, 3, 3)	100	1.7%	1.8%	1.8%	2.0%
(18, 3, 5)	60	0.7%	0.8%	0.8%	0.8%
(18, 3, 15)	65	1.0%	1.8%	1.9%	1.9%
(18, 3, 30)	45	0.9%	1.5%	1.5%	1.5%
(18, 3, 60)	35	1.0%	1.0%	1.5%	1.5%
(32, 5, 3)	55	0.9%	1.1%	1.2%	1.2%
(32, 5, 5)	20	0.9%	0.9%	1.2%	1.2%
(32, 5, 15)	25	0.8%	0.8%	0.8%	0.9%
(32, 5, 30)	22	0.9%	1.2%	1.2%	1.2%
(32, 5, 60)	32	0.8%	1.1%	1.1%	1.1%
(48, 7, 3)	25	0.8%	0.8%	1.1%	1.3%
(48, 7, 5)	15	1.0%	1.1%	1.1%	1.1%
(48, 7, 15)	15	0.6%	0.6%	0.9%	0.9%
(48, 7, 30)	15	0.9%	1.1%	1.1%	1.2%
(48, 7, 60)	15	0.9%	1.4%	1.4%	1.4%
(64, 9, 3)	15	1.0%	1.0%	1.1%	1.3%
(64, 9, 5)	15	0.6%	0.6%	0.7%	0.9%
(64, 9, 15)	15	0.6%	0.6%	0.6%	0.6%
(64, 9, 30)	15	0.8%	0.9%	0.9%	1.2%
(64, 9, 60)	15	1.0%	1.0%	1.1%	1.1%
average stab.		0.89%	1.06%	1.15%	1.22%

Even up to *m=10* we get stability levels under 2%. So *m=4* was OK.

Classical in-sample test for *m=4* and *S* from the earlier tests. All is well.

(N,K)	$P=3$ S	$P=3$ RD	$P=5$ S	$P=5$ RD	$P=15$ S	$P=15$ RD	$P=30$ S	$P=30$ RD	$P=60$ S	$P=60$ RD
$(18,3)$	100	1.7%	60	0.9%	65	0.7%	45	1.1%	35	1.3%
$(32,5)$	55	0.9%	20	0.9%	25	0.8%	22	0.7%	32	0.5%
$(48,7)$	25	0.4%	15	1.1%	15	0.3%	15	0.7%	15	1.0%
$(64,9)$	15	0.4%	15	0.4%	15	0.3%	15	0.4%	15	0.3%

Mester and Bräysy (2007) Active-guided evolution strategies for large-scale capacitated vehicle routing problems. *Computers & Operations Research* 34(10):2964–2975.

Weak out-of-sample test

	$P=3$		$P=5$		$P=15$		$P=30$		$P=60$	
(N,K)	S	SL	S	SL	S	SL	S	SL	S	SL
$(18,3)$	100	1.7%	60	1.2%	65	0.8%	45	1.2%	35	1.3%
$(32,5)$	55	1.3%	20	1.2%	25	1.0%	22	1.1%	32	0.8%
$(48,7)$	25	0.6%	15	1.1%	15	0.6%	15	0.9%	15	1.1%
$(64,9)$	15	0.5%	15	0.6%	15	0.4%	15	0.6%	15	0.8%

So all the stabilities are in place

Kaut and Wallace (2007) Evaluation of scenario generation methods for stochastic programming, *Pacific Journal of Optimization* 3(2) 257-271.

CPU times

T1 – to create the scenarios
T2 – to evaluate the objective function once
S – number of scenarios

	$P=3$			$P=5$			$P=15$			$P=30$			$P=60$		
(N,K)	S	$T1$	$T2$	S	$T1$	$T2$	S	$T1$	$T2$	S	$T1$	$T2$	S	$T1$	$T2$
$(18,3)$	100	160	0.008	60	113	0.008	65	659	0.008	45	2557	0.008	35	21287	0.008
$(32,5)$	55	51	0.014	20	25	0.015	25	342	0.015	22	2273	0.015	32	18097	0.015
$(48,7)$	25	15	0.021	15	20	0.021	15	294	0.021	15	2240	0.022	15	17198	0.022
$(64,9)$	15	9	0.028	15	20	0.028	15	294	0.029	15	2240	0.029	15	17198	0.029

CPU comments

- For the case with P=30 we needed 37 minutes for the 15 scenarios (but they have 180,600 elements) and they handle about 130,000 correlations.
- You need time to create the scenarios – fixed cost
- The number of scenarios shows how much longer it takes per iteration relative to the deterministic case.

Extended tests

- We are in the process of testing alternative VRPs with time windows (hard and soft), fuel consumption (rather than distance), fixed costs for vehicles …

- The results are mostly consistent – about 15 scenarios in the largest cases for stable function evaluations.

- We are also testing Stochastic Shortest Path problems with real data.

Conclusion

- We have a method than can evaluate the objective function for correlated speeds in very high dimensions with very high accuracy. Major issues:
  - An argument that these problems cannot be formulated on a customer-node network, but need the underlying map
  - Proper stability testing to settle the number of scenarios
  - A scenario generation method well suited to VRPs
- The number of scenarios is very modest, implying that solving the stochastic case is not too time consuming.

But can this be extended?

- Major issue: We know which correlations are most important since we know the shape of a solution.
- I think there is hope in some facility location models.
- So let’s see what we can do.

Thank you!