

Routing Redux: 12th DIMACS Implementation Challenge on VRP

[December 2022] Despite COVID-related delays that kept us idling at the depot for over a year, the 12th DIMACS Implementation Challenge on Vehicle Routing Problems (VRPs) launched in the fall of 2021 and concluded in April 2022 with a highly interactive online workshop. The workshop drew to a close with a festive awards ceremony that highlighted a sizable list of contributions. This article recaps some of these contributions, recognizes top-performing teams, and outlines related future plans.

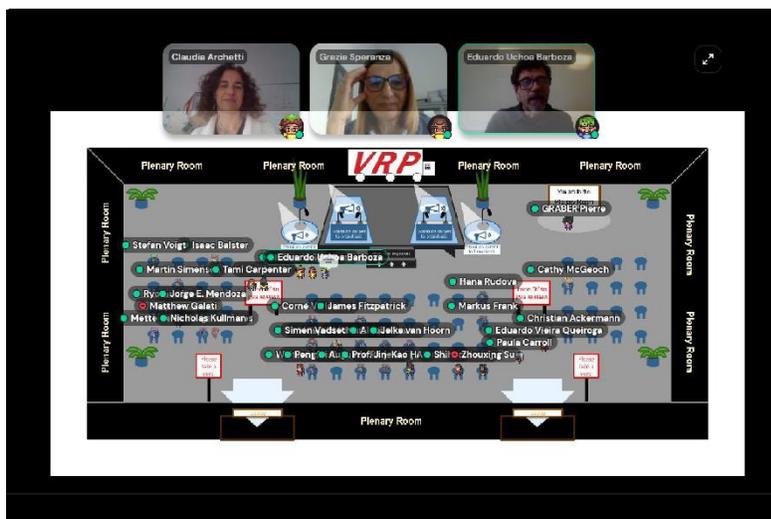


VRPs have been widely studied for decades because they are of both practical relevance and theoretical interest. Designing efficient routes for vehicles performing distribution or service functions translates directly to cost savings, making vehicle routing a topic of great commercial interest and giving rise to an expansive space of problem variations. The Challenge considered seven of these VRP variants:

- 1) Capacitated VRP (CVRP)
- 2) Capacitated VRP with Time Windows (VRPTW)
- 3) Inventory Routing Problem (IRP)
- 4) VRP with Split Deliveries (SDVRP)
- 5) Electric Vehicle Routing (E-VRP)
- 6) Capacitated Arc Routing (CARP)
- 7) Dynamic Ride Hailing (DRH)

These offered a mix of mature VRP variants (CVRP, VRPTW) that have been studied by many groups over multiple decades, adolescent variants (IRP, SDVRP, CARP) whose study has been more recent and/or more limited, and nascent variants (E-VRP, DRH) inspired by new trends in transportation. These

seven VRP variants formed the seven track of the Challenge. Teams competing in the Challenge were invited to enter as many solvers as they liked to compete in as many tracks as desired. Ultimately, 36 teams contributed a total of 47 submissions. Of these, 25 were invited to present in the [Challenge's concluding workshop](#), which was held online April 4-8, 2022. Short papers describing the solvers and videos from workshop presentations are available on this [summary webpage](#).



Organizers Claudia Archetti and Eduardo Uchoa converse with presenter Grazia Speranza in the Plenary Room during the workshop.

The over-arching purpose of an Implementation Challenge is to assess the current state-of-the-art performance of algorithms for a particular problem class. To this end, Challenges often involve curating benchmark instances, developing instance generators, and standardizing formats and formulations for making comparisons. But Challenges are not only about software and data. They are also about fostering interactions and transferring ideas among researchers working on a common class of problems. By all of these measures, the 12th Implementation Challenge was a success.

The Classics: CVRP & VRPTW

The goals of a Challenge differ based on the maturity of the variant under consideration. For well-studied variants like CVRP and VRPTW, the goal was to get a clear sense of the state-of-the-art with respect to the top algorithms and their performance. These tracks, led by Eduardo Uchoa (Universidade Federal Fluminense), proceeded in two phases: 1) a qualifying phase in which competitors ran their solvers on their own computers to solve a common set of test instances; and 2) a final phase in which organizers ran the qualifying solvers on identical DIMACS computers to solve a set of problems unknown to the competitors. In all, eight (of 16) solvers were invited to the final phase for CVRP and seven (of 12) were invited for VRPTW. As a first observation, we note that all of the solvers that advanced to the final round in both the CVRP and VRPTW tracks were based on metaheuristics.

The [final rankings](#) were revealed by Uchoa during a suspenseful awards ceremony marking the end of the workshop. The Alkaid-X team from Huawei Cloud won a tight competition in the CVRP track. Four of the top five solvers in that track (including [Alkaid-X's FHCSolver](#)) were based on the open-source Hybrid Genetic Search (HGS) solver of Challenge co-organizer Thibaut Vidal (Polytechnique Montréal) [V22]. In a closer examination of the results by Uchoa, HGS emerged as the dominant solution approach for all but the largest CVRP instances. For large instances, the ranking changed a bit, with two solvers based on iterated local search joining Alkaid-X in the top three. Interestingly, the Alkaid-X solver uses HGS for smaller instances and iterated local search for the larger instances, resulting in robust performance for all problem sizes. The VRPTW track was won convincingly by the team Wouter & Co from the logistics company ORTEC. [Wouter & Co's solver](#) (called Router) is also based on HGS and plays a central role in a new competition described below.

The Challenge also led to improvements in sets of test instances for these two tracks.

- Two teams—Wouter & Co and Vavavuma! from the University of Bologna—combined to find improving solutions for 160 of the 220 VRPTW instances that were not already optimal in the widely used Homberger and Gehring [HG99] set of 300 instances.
- A set of twelve new CVRP instances based on real data were contributed for use in the Challenge. Six are from ORTEC and based on data from a grocery delivery service in the US, and six are from the Brazilian logistics company Loggi and derived from real data in several large cities in Brazil. These twelve instances were the only ones in the competition to use roadway distances rather than Euclidean distances.
- The twelve new instances are now part of [CVRPLib](#). Also new to CVRPLib

The screenshot shows the CVRPLib website interface. At the top, there are navigation tabs: 'All Instances', 'Plots', 'New Instances', 'CVRP Challenge', 'Updates', 'Posts', 'Links', and 'About'. Below the navigation is a table of benchmarks and instances. The table has columns for Benchmark, Instance, n, K, Q, UB, Opt, and Features. The table lists various benchmark sets like Set A, Set B, Set E, Set F, Set M, Set P, Christofides, Rochat and Taillard, Golden et al., Li et al., Uchoa et al., and Arnold, Gendreau and Sörensen. It also lists DIMACS (2021) instances and VRPTW benchmarks like Solomon (1987) and Homberger and Gehring (1999).

Benchmark	Instance	n	K	Q	UB	Opt	Features
▶ Set A (Augerat, 1995)							🔍 📄 📁
▶ Set B (Augerat, 1995)							🔍 📄 📁
▶ Set E (Christofides and Ellon, 1969)							🔍 📄 📁
▶ Set F (Fisher, 1994)							🔍 📄 📁
▶ Set M (Christofides, Mingozzi and Toth, 1979)							🔍 📄 📁
▶ Set P (Augerat, 1995)							🔍 📄 📁
▶ Christofides, Mingozzi and Toth (1979)							🔍 📄 📁
▶ Rochat and Taillard (1995)							🔍 📄 📁
▶ Golden et al. (1998)							🔍 📄 📁
▶ Li et al. (2005)							🔍 📄 📁
▶ Uchoa et al. (2014)							🔍 📄 📁
▶ Arnold, Gendreau and Sörensen (2017)							🔍 📄 📁
▼ DIMACS (2021)							🔍 📄 📁
	Loggi-n401-k23	400	23	100	336903	no	🔍 📄 📁
	Loggi-n501-k24	500	24	120	177176	no	🔍 📄 📁
	Loggi-n601-k19	600	19	180	113155	no	🔍 📄 📁
	Loggi-n601-k42	600	42	60	347059	no	🔍 📄 📁
	Loggi-n901-k42	900	42	120	245301	no	🔍 📄 📁
	Loggi-n1001-k31	1000	31	180	284356	no	🔍 📄 📁
	ORTEC-n242-k12	241	12	125	123750	no	🔍 📄 📁
	ORTEC-n323-k21	322	21	150	214071	no	🔍 📄 📁
	ORTEC-n405-k18	404	18	150	200985	no	🔍 📄 📁
	ORTEC-n455-k41	454	41	70	292485	no	🔍 📄 📁
	ORTEC-n510-k23	509	23	145	184529	no	🔍 📄 📁
	ORTEC-n701-k64	700	64	80	445543	no	🔍 📄 📁
▶ XML100 (2021)							🔍 📄 📁
VRPTW benchmarks							
Benchmark	Instance	n	K	Q	UB	Opt	Features
▶ Solomon (1987)							🔍 📄 📁
▶ Homberger and Gehring (1999)							🔍 📄 📁

A new set of CVRP instances, DIMACS (2021), was added to CVRPLib for the Challenge. Two classic sets of VRPTW instances were also added.

are the classic VRPTW instances of Solomon [S87] and of Homberger and Gehring [HG99] that were used in the Challenge.

The Adolescents: IRP, SDVRP, CARP

The goals of the Challenge remained the same for the less mature variants but with less need for fine-grained control. This allowed competitors to perform all tests on their own machines, but with time limits scaled to compensate for differences in processor speed.

Claudia Archetti (ESSEC Business School) led the IRP and SDVRP tracks—two of the most challenging in the competition. In each case, the Challenge led to improved best-known solutions for many of the test problems. The Alkaid-X team’s [AlkaidSD solver](#) prevailed in the SDVRP track, beating or matching the competition on 94 of the 95 test instances. The solver is based on an iterated local search metaheuristic for VRSPSD from the literature with several important updates to improve the local search. These include: an adaptation for SDVRP of Vidal’s SWAP* operator [V22] for efficient route improvement; further speedup through use of a matrix cache containing best moves for a neighborhood; and an improved perturbation mechanism. The result is arguably the current state of the art solver for SDVRP.

IRP combines inventory management with vehicle routing to form a difficult optimization problem. Several of the competing solvers implemented “matheuristics” that integrate metaheuristics and mathematical programming techniques to exploit the features of each. Several were also able to substantially improve on the previous best-known solutions for the test instances. The top-performing NTNU AXIOM team of the Norwegian University of Science and Technology combined and improved two of their recent works to deliver an impressive performance. NTNU AXIOM’s [MrOptimal solver](#) is an exact method that uses an efficient matheuristic previously shown to be effective for IRP to warm-start a branch and cut method for IRP.

Although ten teams registered to compete in the CARP track, only one competitor ultimately submitted results, but the results obtained by Stefan Røpke’s [ALNS++ solver](#) were strong. ALNS++ is a brand new implementation based on some of Røpke’s previous work on an adaptive large-neighborhood search (ALNS) metaheuristic. ALNS++ combines ALNS with solving a set covering problem based on the routes generated by ALNS to identify promising areas of the solution space to further explore with ALNS. ALNS++ was able to match, and in a few cases beat, best-known solutions for some of the test instances, and it was competitive with (and sometimes better than) the state-of-the-art unified HGS method of Vidal. ALNS++ also competed in the CVRP track, and it placed fifth in the VRPTW track, making it the only solver to compete in more than two tracks. During the award ceremony, Thibaut Vidal, organizer of the CARP track, noted that arc routing problems have often been viewed as being separate from other types of routing problems. Røpke’s work helps to connect CARP with better-known vehicle routing problems like CVRP and VRPTW.

The Upstarts: E-VRP & Dynamic Ride-hailing

The E-VRP and Ride-hailing tracks were co-led by Nick Kullman (Amazon) and Jorge Mendoza (HEC Montréal). These tracks represent relatively new routing applications that we believe are new to the competition scene. As such, the goal in the Challenge was mainly to spark interest and provide some benchmarks and tools to assist in comparing methods, both during the Challenge and beyond.

Interest in these variants was strong, with 16 teams registered to compete in the E-VRP track and 13 registered in RHP. Nonetheless, barriers to completing the competition were also high precisely because these variants are so new. With few existing resources and codes to build upon, only two teams ultimately submitted results for E-VRP and four for ride-hailing.

The ride-hailing problem, which consists of a central operator controlling a fleet of vehicles that serves randomly arising trip requests, was the only dynamic variant considered in the Challenge. The goal of competitors in this track was to develop a policy for the operator to follow that maximizes profit over a defined planning period. Kullman and Mendoza developed a new open-source python package called pyhailing to support the competition. Pyhailing provides the simulator (an OpenAI Gym environment) with which participants' and their policies interacted during the competition. As requests arrived, competitors determined which requests to serve by which vehicles and how to reposition vehicles to best serve future requests. The winning team, BWOR from the University of Hildesheim, focused on honing its repositioning policy. Given the particular setting simulated in the Challenge, the team opted to combine a relatively simple nearest assignment strategy for assigning requests to vehicles with a more sophisticated repositioning strategy seeking to balance the demand for vehicles in different zones service area with the supply of vehicles. BWOR's [BWOR-L20 solver](#) bested the competition in all four instance categories—small, medium, and large untimed instances, as well as large timed instances.

E-VRP is notoriously difficult, and no team was able to outperform the current state-of-the-art results presented in [FJML22], but the [Redoute solver](#) of a team from University College Dublin offered a promising new approach that was able to match the best-known solutions on smaller instances with a common set of parameters and on some larger instances with tuned parameters. The solver focuses on pruning the search space by judiciously fixing a large fraction of the edge variables using a combination of rule-based approaches and machine learning classification models to predict subsets of edges that can be pruned with minimal effect on solution quality.

Videos of workshop presentations and extended abstracts describing the solvers are accessible on the [Final Papers & Videos](#) page of the Implementation Challenge website.

Related Activities

There are several notable follow-ups on the Challenge. Logistics company ORTEC sponsored a vehicle routing competition called "[EURO Meets NeurIPS](#)" that built on the Challenge and other VRP competitions. It focused on static and dynamic versions of VRPTW, and competitors were required to compete in both tracks. The overarching goal of the competition was to bring together teams working on routing problems from both the operations research and the machine learning communities to bridge gaps and reap benefits of both approaches to further improve solvers. Wouter Kool, the main organizer of the EURO Meets NeurIPS competition, led the team that won the VRPTW competition in the DIMACS Challenge, and the Router solver that won the VRPTW track was open-sourced for use in the new competition. Results were announced by Kool in a workshop associated with the NeurIPS conference in December 2022.

Implementation Challenge organizers arranged for special sections of two different journals to allow authors a choice of publication venue. A section of *Transportation Science* will be devoted to methodological contributions and a section of *INFORMS Journal on Computing* will be devoted to software contributions. Both special sections are open to contributions from both the Implementation Challenge and the EURO Meets NeurIPS competition.

Finally, we are planning to hold an in-person workshop on vehicle routing in May 2023. The themes will expand on those of the Implementation Challenge and emphasize emerging trends, such as solving dynamic vehicle routing problems and combining traditional optimization-based approaches with machine learning. The Challenge was dedicated to David S. Johnson, and the online workshop began with a keynote presentation by Mauricio Resende (Amazon) paying tribute to him. We plan a larger tribute to Johnson at the upcoming workshop.



Mauricio Resende delivers the workshop's opening keynote address honoring David Johnson and his contributions to the theory and implementation of algorithms.

Selected Links:

- Implementation Challenge Website: <http://dimacs.rutgers.edu/programs/challenge/vrp/>
- Challenge results: <http://dimacs.rutgers.edu/programs/challenge/vrp/results/>
- Final Papers & Videos: <http://dimacs.rutgers.edu/programs/challenge/vrp/papers-videos/>
- Challenge Workshop: <http://dimacs.rutgers.edu/events/details?eID=1090>
- EURO Meets NeurIPS: <https://euro-neurips-vrp-2022.challenges.ortec.com/>

References:

- [FJML22] A. Froger, O. Jabali, J.E. Mendoza, G. Laporte. The electric vehicle routing problem with capacitated charging stations, *Transportation Science*, 26(2), 2022.
- [HG99] J. Homberger and H. Gehring. Two evolutionary metaheuristics for the vehicle routing problem with time windows, *INFOR: Information Systems and Operational Research*, 37(3):297–318, 1999.
- [S87] M. M. Solomon. Algorithms for the vehicle routing and scheduling problems with time window constraints, *Operations Research*, 35(2):254–265, 1987.
- [V22] T. Vidal, Hybrid. genetic search for the CVRP: Open-source implementation and SWAP* neighborhood, *Computers & Operations Research*, Volume 140, 2022.