

ODH | CPLEX

Recent Developments and Results

Alkis Vazacopoulos
Robert Ashford
Optimization Direct Inc.
November 2018



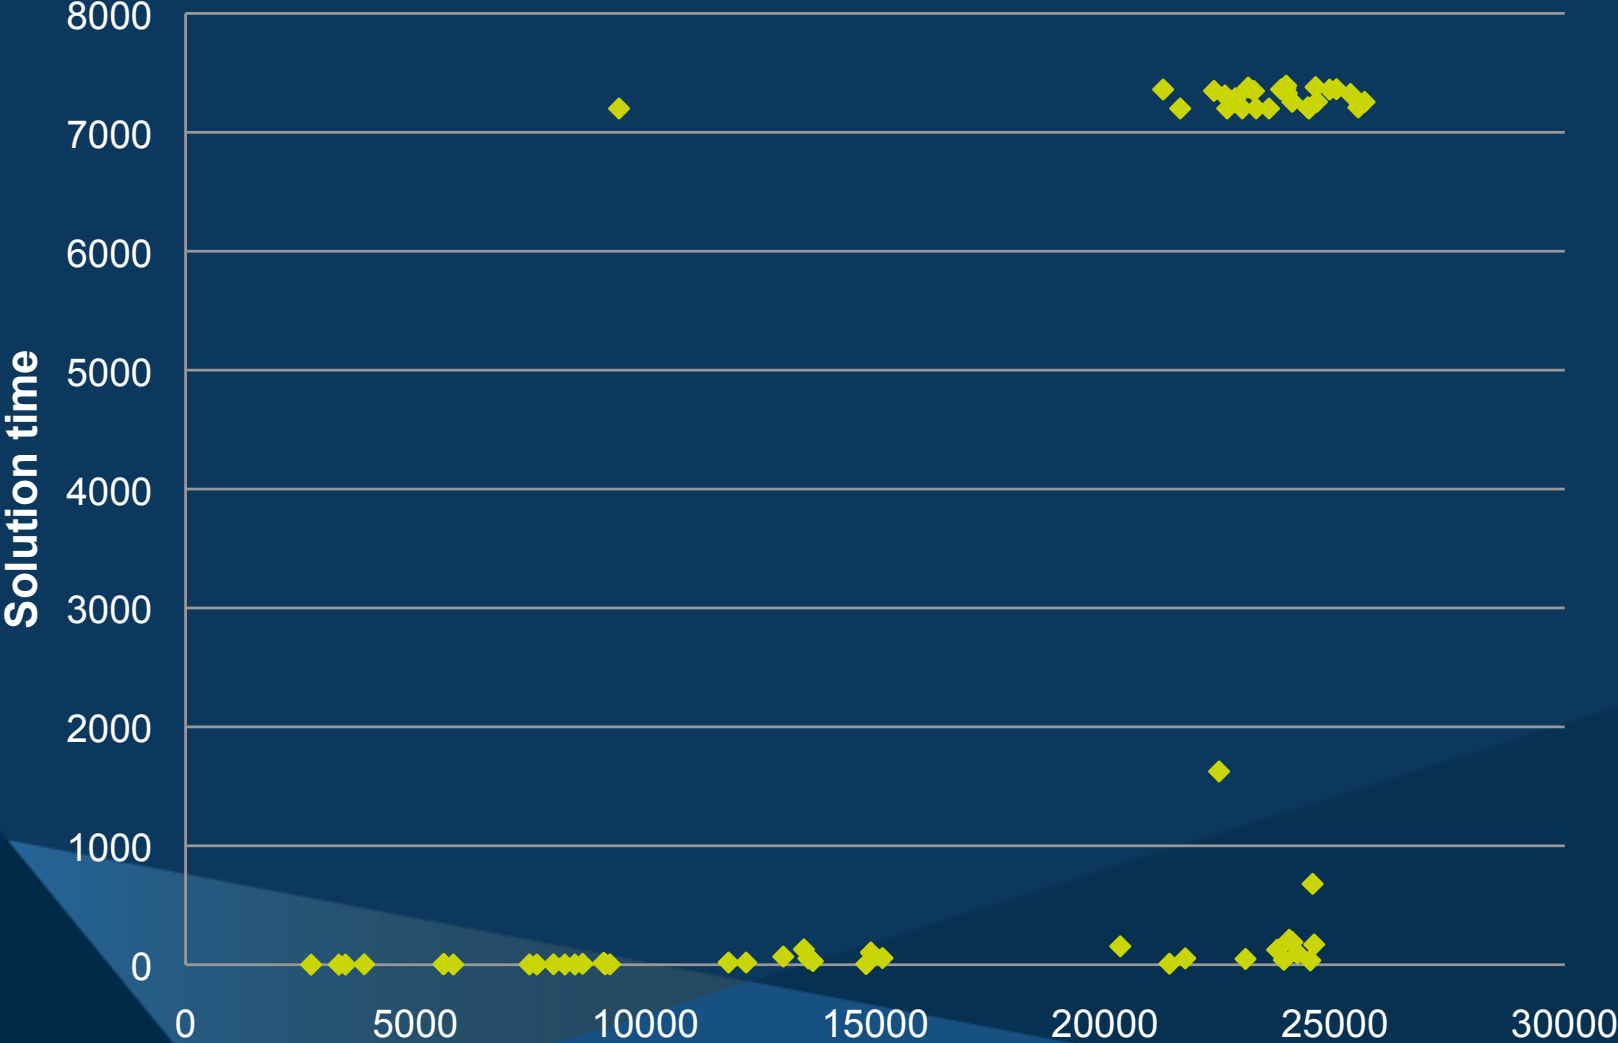
Summary

- Challenges of Large Scale Optimization
- The ODHeuristics approach
 - ODHeuristics Engine
 - ODH|CPLEX Optimizer
- New features in release 4
- Scheduling, supply chain, telecomms and benchmark examples

The Problem: Large Scale Optimization

- Models becoming larger and more complex
- Standard optimization technology stretched/fails
- Super-linear solve time growth often supposed
- *The reality is worse*
 - See how solve time varies with integers after presolve

Solution time vs Size



**OPTIMIZATION
DIRECT**

Integer entities



ODHeuristics: What Is It?

- Tools for
 - handling large and/or difficult MIPs
 - exploits parallel hardware
 - typical server/workstation architecture
 - produces good solutions
 - uses CPLEX for solving sub-models
- ODHeuristics Engine
 - can be used on its own to find solutions
 - But doesn't give optimality guarantee (gap)
- ODH|CPLEX Optimizer
 - CPLEX with the ODHeuristics engine inside
 - Good at getting solutions
 - Gives optimality guarantee

ODHeuristics Engine

- Presented as a software library
 - For embedding into customer applications
 - Call-backs and controls
 - In C, C++/Concert, Java and Python
 - Supports Windows and Linux
- Driver programs are supplied
 - For command line use
 - As examples of calling the library
- Short User Guide (PDF)
- Skeleton scripts for compiling callers and linking



ODHeuristics Engine: How Does it Work?

- Finds an initial solution
 - local search; and/or
 - 'bigM' and 'phase1' methods; or
 - using CPLEX
- Improves its current solution
 - Decomposes original model into sub-models
 - Finds better solution to sub-models (not necessarily optimal)
 - Each ODH thread solves its own set of sub-models
 - Combines the solutions across threads
 - Repeats with fresh decomposition
 - Progressively increases sub-model size
- Decomposition
 - Uses structure inferred from variable names and user-supplied pattern; or
 - Using user call-back; or
 - Automatically inferred from matrix structure
- Deterministic or Opportunistic

What's New in Release 4

- Performance enhancements
 - Improved synchronization
 - Now works well under SMT e.g. 8 threads on 4 core Intel i7
 - Smarter sub-model management
 - Reduced memory fragmentation
- Improved SOS support
- MATLAB support
- Python support (2.7, 3.5 and 3.6)
- Flexible user-specified decomposition via call-back



Examples: Large Scale Scheduling, Supply Chain and Telecomms Models

Model	Application	entities	rows	cols	integers
Easy	Scheduling	314	299288	57804	57804
Mixed	Supply Chain	89177	553715	496455	153183
Medium	Scheduling	314	389560	94200	94200
Difficult	Scheduling	406	371964	149132	149132
Large	Supply Chain	302965	2836736	4892396	1827140
Phase_1	Scheduling	283	421650	155336	154828
Huge	Telecomms	27000	2577916	12944400	12944400



ODH | CPLEX Results

8 Threads on Intel 4 core i7-4790K 4GHz 8hr time limit

	ODH 4.0.6			CPLEX 12.8.0		
	Solution	Time	Gap	Solution	Time	Gap
Easy	96	1 hr 20'	0%	96	6 hrs	0%
Mixed	3180103	15'	1%	3156290	41'	1%
	3153039	8 hrs	0%	3153314	8 hrs	0%
Medium	120	8 hrs	36%	1761	8 hrs	96%
Difficult	985	8 hrs	67%	none	8 hrs	∞
Large	11500971	8 hrs	1%	112784719	8 hrs	90%
Phase_1	31990	8 hrs	5%	none	8 hrs	∞
Huge	401	8 hrs	65%	406	8 hrs	66%

ODH | CPLEX Results

4 Threads on Intel 4 core i7-4790K 4GHz

ODH|CPLEX Optimizer
3.2.0

CPLEX 12.7.1

3 (CPLEX) + 1 (ODH) thrds
(det)

4 threads (deterministic)

	Solution	Time	Gap	Solution	Time	Gap
Easy	96	1 hr 13'	0%	96	2 hr 52'	0%
	3156023	7'	1%	3159264	10'	1%
Mixed			0.05			0.05
	3153082	8 hrs	%	3153147	8 hrs	%
Medium	126	8 hrs	39%	1797	8 hrs	96%
Difficult	1564	8 hrs	79%	none	8 hrs	∞
	1.14820e+			1.23548e+		
	7	8 hrs	1%	7	8 hrs	0.1%
	none	8 hrs	∞	none	8 hrs	∞

ODH|CPLEX Results

12 Threads on Intel 24 core Xeon E5-2690v3 3GHz

ODH|CPLEX Optimizer 3.2.0

CPLEX 12.7.1

9 (CPLEX) + 3 (ODH) thrds
(det)

12 threads (deterministic)

	Solution	Time	Gap	Solution	Time	Gap
Easy	96	4 hr 11'	0%	96	3 h 19'	0%
	3172067	7' 20"	1%	3164927	13' 1"	1%
Mixed	3152986	8 hrs	0.04%	3151448	8 hrs	0%
	120	8 hrs	36%	1760	8 hrs	96%
Difficult	854	8 hrs	62%	none	8 hrs	∞
	1.146618e+	8 hrs	1%	1.414254e+	8 hrs	
Large Phase	1.146618e+	8 hrs	10%	1.414254e+	8 hrs	



More Customer Models: Statistics

Name	Source/Application	rows	cols	nonz	binaries	SOS ints	SOS
Rprs	Modelling System Vendor	126367	107788	431668	23652		
Ship4_14	Modelling System Vendor	732311	719937	3384573	664678		
Ship4_8	Modelling System Vendor	343629	336960	1538418	299656		
20160824	Modelling Consultant	138562	109964	552333	69000		
20160920	Modelling Consultant	108675	81488	415981	54000		
822_p1	Modelling Consultant	395435	543700	2235636	8695	7050	
822_p2	Modelling Consultant	395439	543700	2235640	8695	7050	
824_p1	Modelling Consultant	45840	70697	263418	10944	4119	
Apso374-2	Modelling Consultant	579250	509399	2457594	72202		
Big_p1	Modelling Consultant	866477	2	4788339	87721	1	
Dsj-12p1	Modelling Consultant	40474	44304	177842	9552	4	
IG_2016101	Modelling Consultant						



OPTIMIZATION
DIRECT



More Customer Models: Statistics

Name	Source/Application	rows	cols	nonz	binaries	ints	SOS
				378850	2214218		
JDI	Forestry Mngmt	1218099	7	0	693064		
				332839	1540679		
JDI_Aug1	Forestry Mngmt	884565	9	3	73997		
				807288	3778687	501784	
Ilim_Model	Forestry Mngmt	6383910	5	2	9	8329	
pulp	Field Installation Mngmt	33464	118451	1579189	118450		
				3324451			
M297C	Healthcare	20496	1625	2		1620	
				3324451			
M305B	Healthcare	20496	1625	2		1620	
				3324451			
Shim	Healthcare	20496	1625	2		1620	
				3324451			
M297C_mat	Healthcare	20496	1625	2		1620	
	Modelling System			107323			
Vendor		497290	0	3330837	8691		7642
OPTIMIZATION	Wireless Ind Strat						
DIRECT							
F880ba86	Mngmt	745732	30332	1518188	30332		



OPTIMIZATION
DIRECT



More Customer Models: Results

8 Threads	ODH CPLEX 4.0.6		CPLEX12.8.0	
Model	Gap	Time	Gap	Time
Rprs	0%	222	0%	714
Ship4_14	7%	7201	17%	7201
Ship4_8	5%	7201	6%	7201
20160824	60%	7200	91%	7200
20160920	21%	7200	39%	7200
822_p1	0%	145	0%	397
822_p2	0%	113	0%	191
824_p1	0%	3	0%	12
Apso374-2	1%	7203	22%	7211
Big_p1	0%	324	0%	212
Dsj-12Jan_p1	3%	7203	6%	7235
IG_20161019	11%	7200	7%	7201

More Customer Models: Results

8 Threads	ODH CPLEX 4.0.6		CPLEX12.8.0	
Model	Gap	Time	Gap	Time
JDI	66%	7206	1600%	7213
JDI_Aug1	25%	7203	inf	7203
Ilim_Model	0%	528	0%	220
pulp	0%	279	0%	646
M297C	19%	7203	42%	7329
M305B	33%	7202	39%	7335
Shim	18%	7205	34%	7502
M297C_mat	14%	7204	47%	7714
ict1	0%	3975	0%	4034
f880ba86-b266	10%	7200	26%	7200
P_201x201	0.10%	7200	0.12%	7201
P_203x203	0.39%	7201	8%	7200

More Customer Models: Results

24 Threads	ODH CPLEX 4.0.6		CPLEX12.8.0	
Model	Gap	Time	Gap	Time
Rprs	0%	193	0%	246
Ship4_14	6%	7203	8%	7204
Ship4_8	5%	7202	6%	7202
20160824	13%	7201	92%	7200
20160920	14%	7201	41%	7201
822_p1	0%	61	0%	450
822_p2	0%	92	0%	323
824_p1	0%	4	0%	11
Apso374-2	1%	7207	1%	7206
Big_p1	0%	612	0%	292
Dsj-12Jan_p1	7%	7222	4%	7226
IG_20161019	4%	7201	9%	7201



More Customer Models: Results

24 Threads	ODH CPLEX 4.0.6		CPLEX12.8.0	
Model	Gap	Time	Gap	Time
JDI	66%	7206	2%	7207
JDI_Aug1	24%	7211	inf	7207
Ilim_Model	0%	423	0%	360
pulp	0%	331	0%	630
M297C	23%	7212	40%	7209
M305B	18%	7207	41%	7209
Shim	19%	7212	46%	7209
M297C_mat	18%	7210	43%	7207
ict1	0%	4722	0%	5946
f880ba86-b266	7%	7201	20%	7201
P_201x201	0.09%	7208	0.11%	7210
P_203x203	0.26%	7200	7%	7201

More Customer Models: Results

24 Threads	CPLEX 12.7.1		ODH CPLEX 3.2.0	
Model	Gap	Time	Gap	Time
Rprs	0%	516	0%	273
Ship4_14	8.3%	7200	6.8%	7200
Ship4_8	5.3%	7200	5.3%	7200
20160824	92.8%	7200	32.7%	7200
20160920	22.2%	7200	9.9%	7200
822_p1	0%	479	0%	94
822_p2	0%	276	0%	113
824_p1	0%	9	0%	2
Apso374-2	21.9%	7200	0.5%	7200
Big_p1	0%	276	0%	350
Dsj-12Jan_p1	4.9%	7200	4.4%	7200
Gom_p1	0%	0.2	0%	0.4
IG_20161019	11.1%	7200	5.9%	7200



More Customer Models: Results

Summary

2 hour time limit

	ODH CPLEX 4.0.6	CPLEX 12.8.0
	Gap	Gap
Average 8 threads	12%	87%
Average 24 threads	9%	19%

ODH | CPLEX Benchmark Results

Mittelmann's 'Solvable' MIPLIB2010 subset

Intel Xeon Gold 6138, 40 cores, 256GB, 2.00GHz, Linux

Current solvable MIPLIB2010 subset less numerically unstable models

Tests run by Hans **NOT** Optimization Direct

40 Threads, 2 hour time limit

Look at relative performance models taking at least τ minutes

τ	CPLEX 12.8.0 takes at least τ			ODH CPLEX takes at least τ		
	# models	s.g.m.	a.m.	# models	s.g.m.	a.m.
0	220	-1.6%	+7.4%	220	-1.6%	+7.4%
5	37	+19%	+12%	43	-10%	+4.7%
10	27	+21%	+13%	29	-17%	+4.1%

ODH | CPLEX Benchmark Results

Mittelmann's 'Slightly Pathological' Models

"This benchmark is not giving a representative impression of the relative performance of the codes"

Intel Xeon X5680 (32GB, Linux, 64 bits, 2*6 cores)

Tests run by Hans **NOT** Optimization Direct

12 Threads, 3 hour time limit

Look at relative performance models taking at least τ minutes

τ	CPLEX 12.8.0 takes at least τ			ODH CPLEX takes at least τ		
	# models	s.g.m.	a.m.	# models	s.g.m.	a.m.
0	40	+12.6%	+7.5%	40	+12.6%	+7.5%
5	30	+15%	+7.5%	26	+1.9%	+6.1%
10	23	+13%	+7.8%	23	+0.6	+5.8%

ODH | CPLEX Performance

- Relatively better for hard-to-solve models
 - Not necessarily large
 - Can get benefit on small models
- Relatively better when run with more threads
 - Effective with 8 threads on 4 core Intel i-7 using SMT
 - Hardware performance bottleneck currently bus speed



Conclusions

- Customers now want to solve larger and large models
- Hard size barriers to solve (to optimality) or even to getting a solution at all
- ODHeuristics can find good solutions
 - Useful on small(er) models too
- ODH | CPLEX can provide solutions of proven optimality quality
- Parallel solution methods best way of exploiting modern hardware (although limited by memory bus speeds)

Benchmarking and Evaluation

- If you think that ODHeuristics and/or ODH | CPLEX might work for you:
- send us your difficult matrices and we will send you the results
- request an evaluation copy



**OPTIMIZATION
DIRECT**



Thanks for listening

Robert Ashford

rwa@optimizationdirect.com

www.optimizationdirect.com



Why Is There a Problem?


- Many models now solved routinely which would have been impossible (“unsolvable”) a few years ago
- **BUT:** have super-linear growth of solving effort as model size/complexity increases
- **AND:** customer models keep getting larger
 - More and better data are available (“Big Data”)
 - Globalized business has larger and more complex supply chain
 - Optimization expanding into new areas, especially scheduling
 - Detailed models easier to sell to Mngmt and end-users

The Curse of Dimensionality: Size Matters

- Super-linear solve time growth often supposed
- *The reality is worse*
- Few data sets available to support this
- Look at randomly selected sub-models of two scheduling models
 - Simple basic model
 - More complex model with additional entity types
 - Two hour time limit on each solve
 - 8 threads on 4 core hyperthreaded Intel i7-4790K
- See how solve time varies with integers after presolve



Simple model



**OPTIMIZATION
DIRECT**

Numer of ineteger entities



Parallelization and Hardware

- CPU clock speed is not improving
 - Power consumed (and heat generated) $\sim \text{speed}^2$
- Memory speed is not improving (much)
- Can get wider registers (vectorization)
 - of limited use in sparse optimization
- Multi-core (processor) machines
 - Can fit more processors onto a single chip
 - 24 cores now on inexpensive servers
 - Exploited by ODHeuristics and CPLEX multi-threading
 - Cannot always use full processor capability if using many cores
 - Performance limited by bus speed ($\sim 20\text{MB/sec}$)

Further Examples: Large Scale Scheduling and Supply Chain Models

Model	entities	rows	cols	integers
Mixed	89177	553715	496455	153183
Phase1	283	421650	155336	154828



Heuristic Results

8 Threads on Intel 24 core Xeon E5-2690v3 3GHz

ODHeuristics 2.14

CPLEX 12.6.3

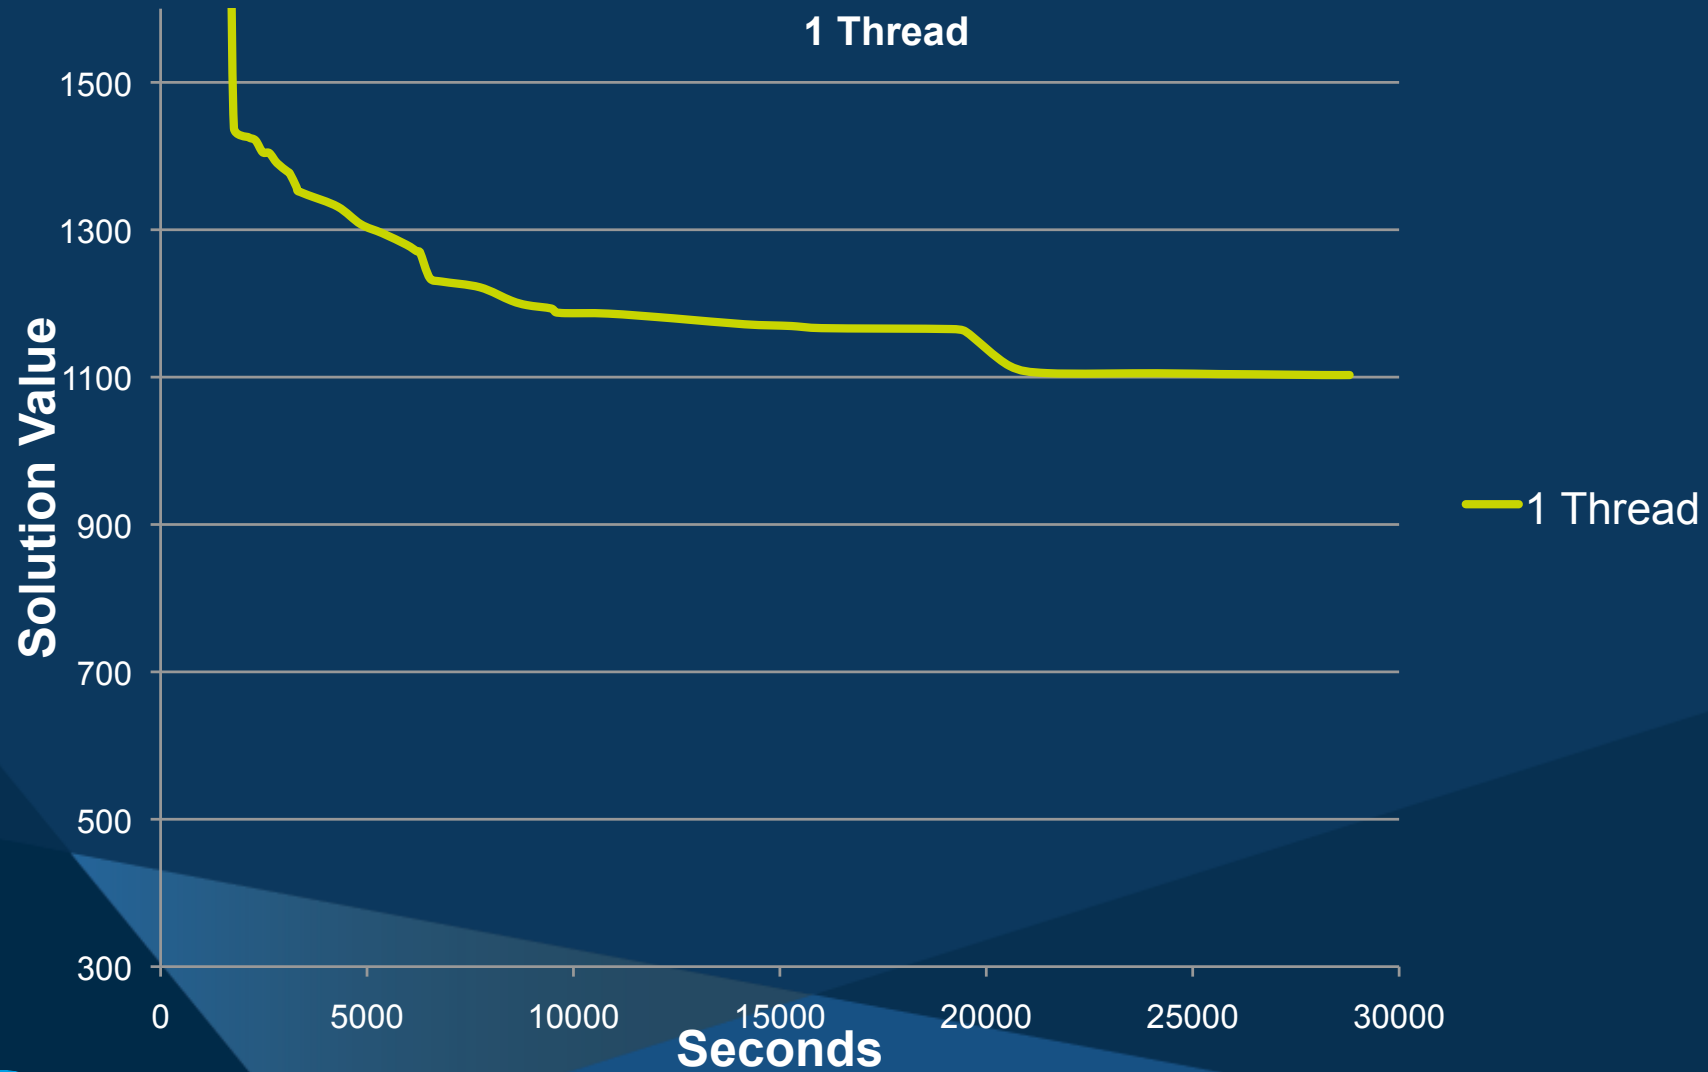
	Solution	Time	Gap	Solution	Time	Gap
Easy	96	4 mins	0%	96	4 ½ hrs	0%
Medium	113	8 hrs	≤ 13%	1161	3 ½ days	93%
Difficult	773.8	8 hrs	≤ 56%	106682	3 ½ days	99.7%
Large	1.149E+7	8 hrs	≤ 1%	1.420E+7	3 days	20%
Huge	370	8 hrs	≤ 61%	412	3 days	66%




**OPTIMIZATION
DIRECT**




Difficult Model Heuristic Behavior



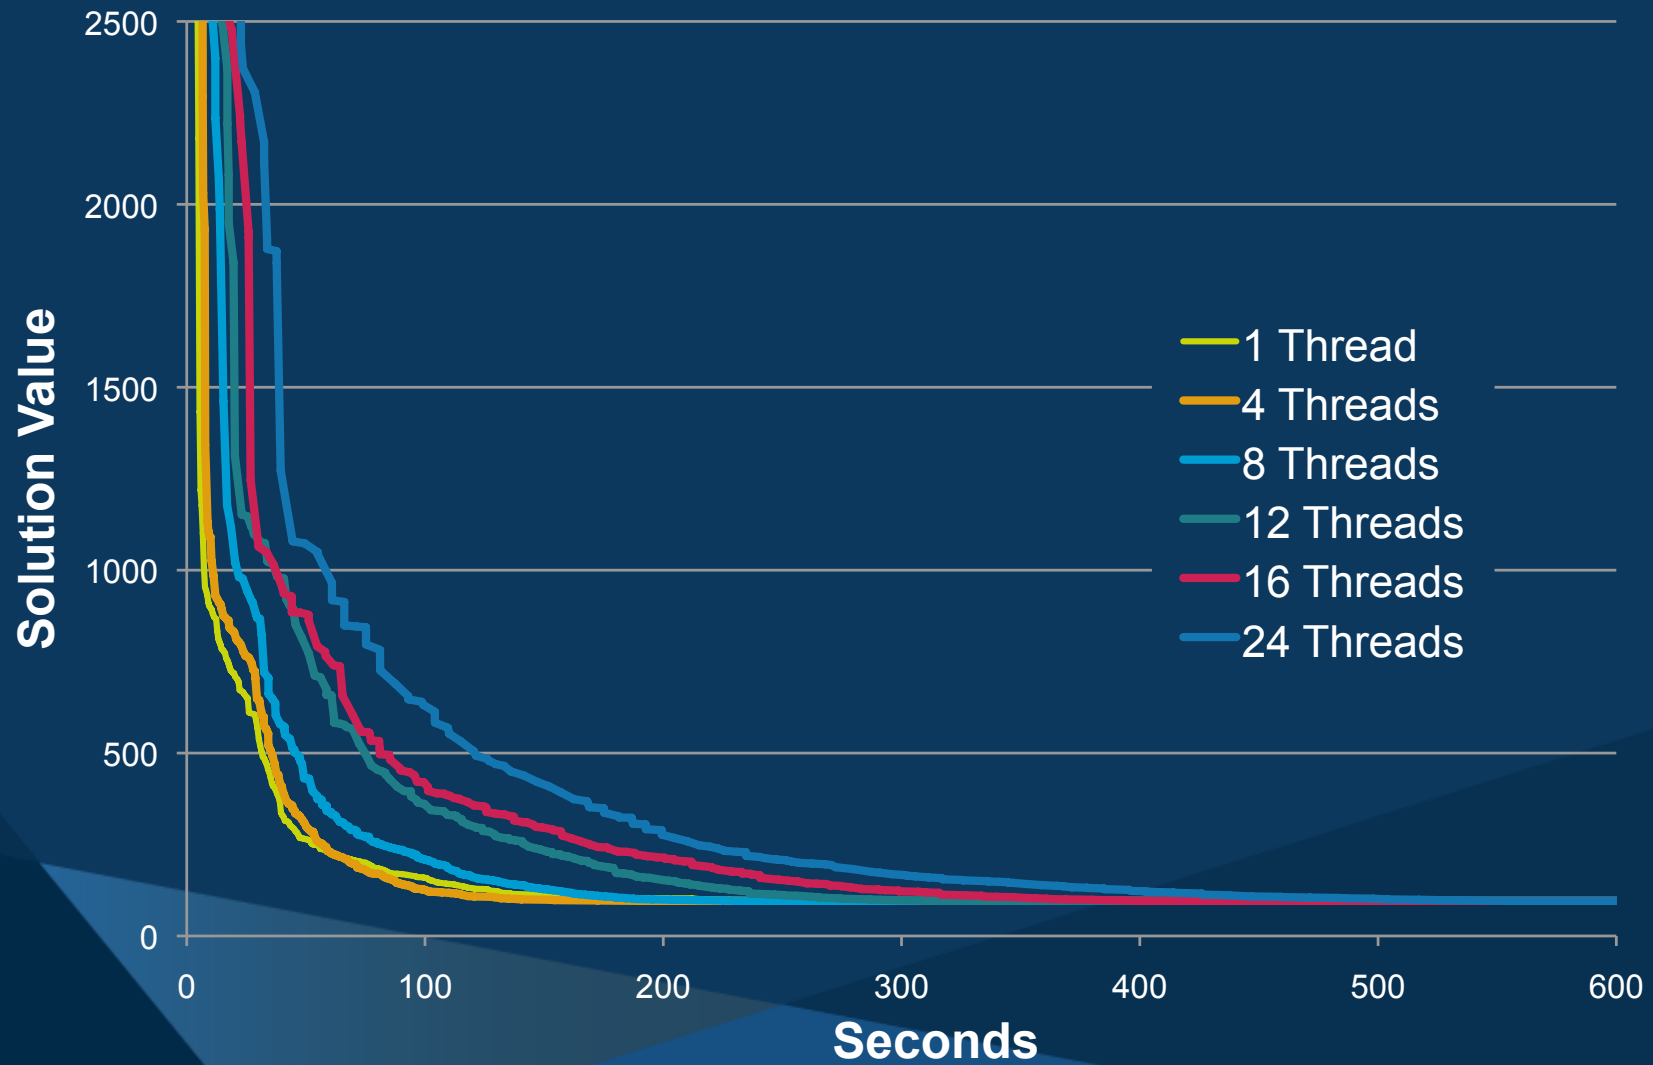
Difficult Model Heuristic Behavior



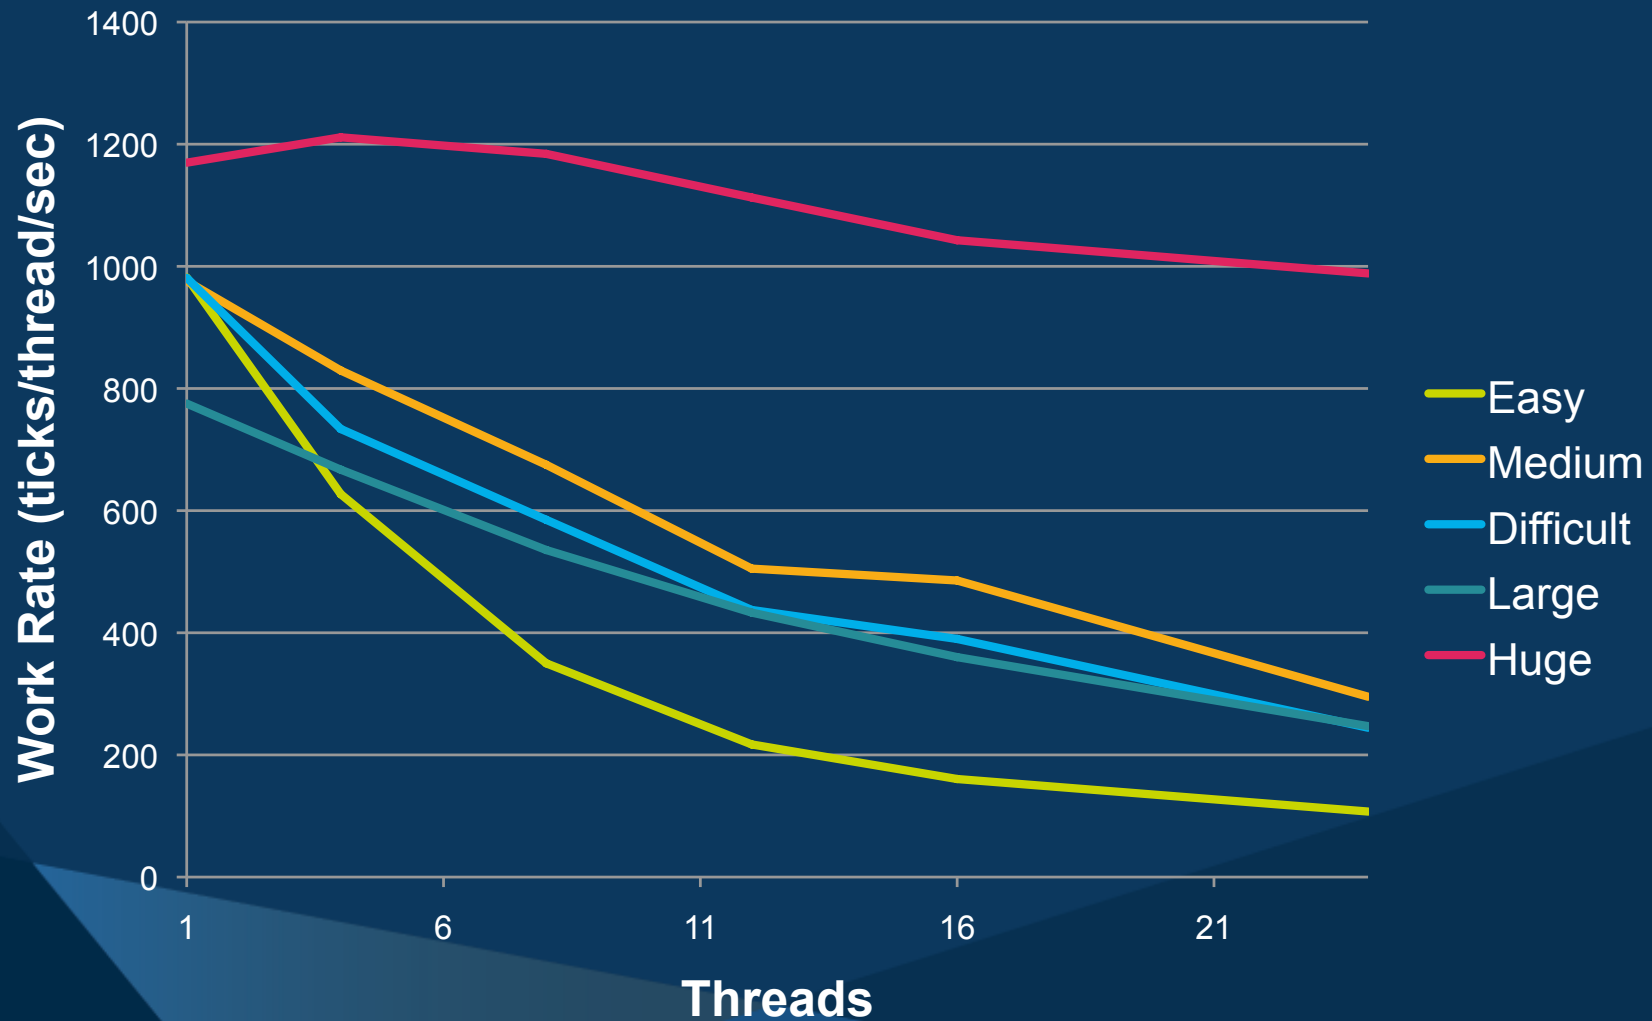
Easy Model Heuristic Behavior



Easy Model Heuristic Behavior



Work Rate



**OPTIMIZATION
DIRECT**



ODH | CPLEX Optimizer

- ODH Engine run underneath CPLEX
- Co-running continuously in separate (sets of) threads
- CPLEX initiates run of ODH
- Either can terminate and stop the other
 - time limit, gap, etc.
- Information exchanged between CPLEX and ODH Engine
- Deterministic or opportunistic modes


ODH | CPLEX Design

- Presented as a subroutine library like CPLEX itself
- All CPLEX API calls are available including
 - Call-backs
 - Parameter interrogation and setting
- Just a different optimize call
- Behaves like CPLEX if ODH doesn't find anything useful (as much as technically possible)
- Supports ODH engine controls and call-back
- Supports C++/Concert and Java



ODH | CPLEX Optimizer Advantages

- Potentially faster than either alone
- Good solutions from ODH accelerate CPLEX, e.g:
 - reduced cost fixing
 - **tree pruning**
 - helps CPLEX heuristics
 - saves effort
 - provides better starting information (e.g. RINS)
- Good solutions from CPLEX accelerate ODH Engine
 - focus ODH's attention on different parts of
- Use more cores than ODH heuristics stand-



ODH | CPLEX Optimizer Challenges

- Synchronization overheads in deterministic mode
 - For CPLEX up to 20% but usually less than 5%
- CPLEX solution paths are changed by just the presence of communication channel with ODH (even if no information is exchanged)
- Needs more memory than CPLEX or ODH alone
 - Twice as much
- Needs more processors/cores
 - 8+ cores to run well, 24+ cores ideal
 - Hyperthreaded (SMT) core splitting not a good idea
- Increases demand on memory bus
 - But less than if all threads were committed to ODH alone

ODH | CPLEX Results

- Usually very beneficial on hard user models
 - Especially scheduling
- Can be of benefit in general
 - Relatively better on harder-to-solve models
 - Relatively better with more cores (>4)
- Tested on customer models
- Tested on subsets of MIPLIB2010 used by Mittelmann

ODH | CPLEX Customer Models: Statistics

Model	Rows	Cols	Nonzeros	Binaries	Integers	SOSs
abaco	181170	306165	983149	5566	542	
allocam	1709	22888	63965	22822		
b19374	182588	104038	1642666	100601		
bp_case3	126128	269091	907017	9214		
case299	169807	142764	2439949	9415		
coal0	10977	11924	83999		2842	
cstscm2	186000	235715	758216	4947	609	
dpa	65417	126402	573546	59800		
exp01	1602710	1574843	7437368	1464408		
exp02	2628020	2617643	12652328	2507208		
fra2	3340550	160413	14930048	157488	1050	
fra21	31225	182113	430866	157488	1050	
fra22	31592	184104	435907	159112	1132	



ODH | CPLEX Customer Models: Statistics

Model	Rows	Cols	Nonzeros	Binaries	Integers	SOSs
fra23	31569	183357	434422	158365	1132	
fra2r	31230	182113	432486	157488	1050	
gcbp	10569	8311	74428	897	210	
genmod3	38978	41856	99826	16368		
harvest1	795	17419	418553	17171		
kat1	497290	1073230	3330837	8691		7642
kat2	679266	1538689	5076733	15628		13404
maletraits	35049	951	76080	890		
mbm	22100	30100	450980	30000		
mem_cg	35971	72578	192419	31972		
nuc	1282	3729	14699	3729		
nvsp01	249368	732735	3293146	730892	1634	
para1	1167600	1130706	4381976	14319		



ODH | CPLEX Customer Models: Statistics

Model	Rows	Cols	Nonzeros	Binaries	Integers	SOSs
pgom	76225	265681	954365	9184	1136	
prakcr	10127	31278	237321	30888		
producao	9410	7488	25263	6675		
radiumfoam	76349	49322	476729	32714		
rjune06	243857	362099	1191709	60814		
scd	349601	333122	1538512	1560		
sp3	813225	675862	2953834	25518	14466	
spp1	442204	444054	1640953	132731		
uni01	216288	303133	930668	6007	573	
utopia	970500	957951	3817642	470		
woods	315763	153037	2537381	82723		

ODH | CPLEX Customer Models: Statistics

Summary

Set of 37 hard models chosen by customer

ODH|CPLEX not specifically designed or tuned for them

Model	Rows	Cols	Nonzeros	Binaries	Integers	SOSs
Max	3340550	2617643	14930048	2507208	14466	13404
Min	795	951	14699	470	210	7642
Average	402230	392889	2105946	173160	2110	10523

ODH | CPLEX Customer Models: Results

Model	CPLEX 12.7.0		ODH CPLEX 3.1.8	
	Gap	Time	Gap	Time
abaco	0.00%	4988	0.00%	3103
allocam	0.00%	6307	0.00%	8026
b19374	0.00%	452	0.05%	10800
bp_case3	0.56%	10800	0.60%	10800
case299	2.75%	10800	2.05%	10800
coal0	23.85%	10800	22.50%	10800
cstscm2	0.06%	10800	0.06%	10800
dpa	0.02%	10800	0.02%	10800
exp01	1.18%	10800	1.57%	10800
exp02	0.06%	10800	0.07%	10800
fra11	0.19%	10800	0.19%	10800
fra21	0.01%	10800	0.00%	344
fra22	0.07%	10800	0.08%	10800



ODH | CPLEX Customer Models: Results

Model	CPLEX 12.7.0		ODH CPLEX 3.1.8	
	Gap	Time	Gap	Time
fra23	1.20%	10800	1.70%	10800
fra2r	0.33%	10800	0.44%	10800
gcbp	68.64%	10800	29.62%	10800
genmod3	∞	10800	∞	10800
harvest1	69.59%	10800	46.46%	10800
kat1	51.52%	10800	22.44%	10800
kat2	89.52%	10800	21.14%	10800
maletraits	4.75%	10800	4.75%	10800
mbm	1.94%	10800	1.44%	10800
mem_cg	11.91%	10800	9.86%	10800
nuc	13.35%	10800	8.64%	10800
nvsp01	0.02%	10800	0.02%	10800
para1	2.03%	10800	2.16%	10800



ODH | CPLEX Customer Models: Results

Model	CPLEX 12.7.0		ODH CPLEX 3.1.8	
	Gap	Time	Gap	Time
pgom	0.60%	10800	0.60%	10800
prakcr	0.00%	4439	0.00%	2327
producao	11.27%	10800	6.31%	10800
radiumfoam	0.02%	10800	0.03%	10800
rjune06	0.07%	10800	0.55%	10800
scd	9.81%	10800	7.16%	10800
sp3	10.05%	10800	3.18%	10800
spp1	0.21%	10800	0.23%	10800
uni01	0.00%	2543	0.00%	3198
uria1	0.21%	10800	0.25%	10800
utopia	0.84%	10800	0.72%	10800
woods	∞	10800	∞	10800

ODH | CPLEX Customer Models: Results

Summary

Tests run by customer on 12 threads on 6 core Intel Xeon win64
3 hour time limit

	CPLEX 12.7.0		ODH CPLEX 3.1.8	
	Gap	Time	Gap	Time
Average	10.46%	9026	5.41%	8809

ODH | CPLEX Developments Underway

- Support for CPLEX 12.8
 - Won't harm CPLEX search strategies by increasing tick counts
- New automatic matrix decomposition method
 - Key to effective use of the technology if users do not want to specify decomposition
- Direct Python support
 - Python users currently use C API