On Some Network-Structured Mixed-Integer NonLinear Problems with Applications to IP Routing

(MINLP meets computer networks)

A. Frangioni

with L. Galli M.G. Scutellà G. Stea

Dipartimento di Informatica, Università di Pisa
Dipartimento di Ingegneria dell'Informazione, Università di Pisa

COST Workshop on Mixed Integer Nonlinear Programming
Paris, October 1st, 2013

Outline

Motivation

- Trends in computer networks:
  - high-bandwidth applications
  - stringent Quality of Service (QoS) guarantees
- Packet-switching (IP, Ethernet) now dominant
- Issue: real-time guarantees (e.g. controlled end-to-end delay)
- Relevant in many WAN/LAN settings:
  - industrial control systems
  - remote sensing and surveillance systems
  - live Internet Protocol Television
  - IP Telephony
  - online gaming/MMORPGs
- Critical in embedded systems (automative, avionics, ...)
- Packets, not circuits $\Longrightarrow$ requires traffic engineering

Delay Constrained Routing Problem

- IP Network ≡ directed graph $G = (N, A)$ ($n = |N|$, $m = |A|$), MTU L
- Set of flows K: origin/destination (s^k, d^k), deadline δ^k
- Leaky-bucket traffic shaper ≡ Arrival curve $\mathcal{A}$ ≡ parameters σ^k (burst) and ρ^k (rate) ($\mathcal{A}(t) = \sigma + t\rho$)
- $(i, j) \in A$: link delay l_{ij}, link speed w_{ij}, reservable capacity c^k_{ij} ($\leq w_{ij}$)
- $i \in N$: node delay n_i
- Linear capacity reservation cost f_{ij} (often $= 1$ ≡ Equal Cost (EC))

Delay Constrained Routing Problem (DCR)

Compute paths and reserve resources on arcs at minimum cost such that the maximum delay of each flow is $\leq$ deadline

- Single-Flow Single-Path (SFSP) DCR: $|K| = 1$, unsplittable flow

Worst-case delay modeling

- One new flow to enter (drop superscripts, r_{ij}^k = existing flows, fixed)
- Worst-case delay of the flow depends on several factors:
  1. the selected s-d path P in G;
  2. the reserved rate (capacity) $r_{ij} \in [0, c_{ij}]$ for each $(i, j) \in P$
  3. the details of the scheduling protocol (requires network calculus)
- Necessary assumption for finite delay: $r_{ij} \geq \rho$ for each $(i, j) \in P$ ($\rho \equiv$ rate $\equiv$ "steady-state" flow demand in classical flow models)
- General formula (already nonlinear!):

$$\frac{\sigma}{\min\{\, r_{ij} \;:\; (i,j) \in P \}} + \sum_{(i,j) \in P} \left(\theta_{ij} + l_{ij} + n_i \right) \qquad (1)$$

  where $\theta_{ij} \equiv$ protocol-specific arc delay (also nonlinear!)
- (1) convex and SOCP-representable if θ_{ij} is

Worst-case delay modeling (cont.)

- Exact formula for θ_{ij} depends on the scheduling protocol:

$$\theta_{ij} = \frac{L}{r_{ij}} + \frac{L}{w_{ij}} \qquad \text{Strictly Rate-Proportional} \quad (2)$$

$$\theta_{ij} = \frac{L}{r_{ij}} + |P(i,j)|\frac{L}{w_{ij}} \qquad \text{Weakly Rate-Proportional} \quad (3)$$

$$\theta_{ij} = \frac{L}{w_{ij}} \frac{w_{ij} - r_{ij}}{\min\{\, r_{ij} \,,\, r_{ij}^{min} \,\}} + \frac{L}{r_{ij}} + |P(i,j)|\frac{L}{w_{ij}} \qquad \text{Frame-Based} \quad (4)$$

$P(i,j)$ = set of paths passing through (i,j) excluding new one

$r_{ij}^{min} = \min\{\, r_{ij}^{k} \;:\; q^{k} \in P(i,j) \,\}$ ($\Longrightarrow$ SRP $\lesssim$ WRP $\leq$ FB)

- (2) flow-independent, convex, SOCP-representable
- (3) $\approx$ (2) but not flow-independent
- (4) (surprisingly) also convex but only for SFSP, less trivial
- (3) and (4) not flow-independent $\Longrightarrow$ have admission control issue

A SOCP model for SFSP-DCR [with SRP]

- Path binary variables x_{ij}, reserve continuous variables r_{ij}

$$\min \textstyle\sum_{(i,j)\in A} f_{ij} r_{ij} \tag{5}$$

$$\sum_{(j,i)\in BS(i)} x_{ji} - \sum_{(i,j)\in FS(i)} x_{ij} = \begin{cases} -1 & \text{if } i = s \\ 1 & \text{if } i = d \\ 0 & \text{otherwise} \end{cases} \qquad i \in N \tag{6}$$

$$0 \leq r_{ij} \leq c_{ij} x_{ij} \qquad (i,j) \in A \tag{7}$$

$$\rho \leq r_{min} \leq r_{ij} + c_{max}(1 - x_{ij}) \qquad (i,j) \in A \tag{8}$$

$$t + \sum_{(i,j)\in A} \left(\theta_{ij} + \left(\frac{L}{w_{ij}} + l_{ij} + n_i \right) x_{ij} \right) \leq \delta \tag{9}$$

$$t\, r_{min} \geq \sigma \quad , \quad t \geq 0 \tag{10}$$

$$x_{ij} \in \{0,1\} \quad , \quad r_{ij} \in \mathbb{R} \qquad (i,j) \in A$$

- (10) rotated SOCP constraint $\equiv t \geq \sigma / r_{min}$ (since $t \geq 0$)
- Issue: how to write "$x_{ij} = 1 \Longrightarrow \theta_{ij} \geq L / r_{ij}$, $x_{ij} = 1 \Longrightarrow \theta_{ij} = 0$"

Solution I: big-M

- Issue: can't use $r_{ij}\,\theta_{ij} \geq L$ for that $\Longrightarrow \theta_{ij} > 0$ always
- Solution: two extra sets of variables s_{ij} and r'_{ij}

$$0 \leq \theta_{ij} \leq M x_{ij}$$

$$\theta_{ij} \geq s_{ij} - M(1 - x_{ij})$$

$$s_{ij}\, r'_{ij} \geq L \quad , \quad s_{ij} \geq 0$$

$$0 \leq r'_{ij} \leq r_{ij} + M(1 - x_{ij})$$

- $\theta_{ij} \geq s_{ij}$ if $x_{ij} = 1$, while θ_{ij} and s_{ij} are "free" if $x_{ij} = 0$
- $r'_{ij} \leq r_{ij}$ if $x_{ij} = 1$, while r'_{ij} and r_{ij} are "free" if $x_{ij} = 0$
- $s_{ij} \geq L / r'_{ij} \Longrightarrow \theta_{ij} \geq s_{ij} \geq L / r'_{ij} \geq L / r_{ij}$ if $x_{ij} = 1$
- $M = \max(\sqrt{L}\,,\ L/\rho)$ suffices, still it's big-M: can we do better?

Solution II: Perspective Reformulation

- General Perspective Reformulation: $f : \mathbb{R}^q \to \mathbb{R}$ convex, two sets

$$\mathcal{P}_0 = \{\, 0 \,\} \quad , \quad \mathcal{P}_1 = \{\, v \in \mathbb{R}^q \,:\, l \le v \le u \,,\, f(v) \le 0 \,\}$$

  the best possible convex approximation of their (nonconvex) union is

$$conv(\, \mathcal{P}_0 \cup \mathcal{P}_1 \,) = \Big\{\, v \,:\, \lambda l \le v \le \lambda u \,,\, \lambda f(v/\lambda) \le 0 \,,\, \lambda \in [0,1] \,\Big\}$$

- Application: after a little tedious algebra

$$\rho x_{ij} \le r_{ij} \le c_{ij} x_{ij} \quad , \quad 0 \le \theta_{ij} \le (L/\rho) x_{ij} \quad , \quad \theta_{ij} r_{ij} \ge L x_{ij}^2$$

  (now θ_{ij} can be 0 when $x_{ij} = 0$, x^2/r convex for $r > 0$)

- original variables + a(nother rotated) SOCP constraint

- Looks much better: is it?

Instances

- Real-world IP network topologies (`GARR`, `SNDlib`, `TopoZOO`)
- Realistic random topologies (`Waxman` model)
- Equal (reservation) Costs $f_{ij} = 1$
- `FNSS` tool for realistic traffic matrices ($\mu(T) = 0.8$ Gbps and $\sigma^2(T) = 0.05$) and link-capacity assignment ($1, 10, 40$ Gbps)
- `DCR-generator` for the remaining network parameters ($L = 1500$, $n_i = l_{ij} = L/w_{ij}$, $\sigma = 3L$)
- Distributed at `http://www.di.unipi.it/optimize/Data/MMCF.html#UMMCF`
- Experiments with "unloaded networks", but "loaded" case analogous

SOCP models – Cplex

	Cplex P				Cplex bM			
	avg		max		avg		max	
	t	n	t	n	t	n	t	n
abilene	0.011	0.000	0.03	0	0.02	0.03	0.09	1
atlanta	0.015	0.044	0.18	1	0.03	0.07	0.17	1
cost266	0.015	0.017	0.06	1	0.05	0.03	0.26	1
dfn-bwin	0.012	0.000	0.03	0	0.05	0.02	0.11	1
dfn-gwin	0.020	0.151	0.10	1	0.05	0.00	0.16	0
di-yuan	0.051	1.190	0.34	18	0.11	1.36	0.62	31
france	0.014	0.000	0.05	0	0.04	0.02	0.16	1
geant	0.011	0.016	0.06	1	0.03	0.03	0.19	1
germany50	0.024	0.025	0.10	1	0.09	0.06	0.70	1
giul39	0.245	0.547	0.99	13	1.27	15.33	6.68	610
india35	0.021	0.036	0.27	1	0.08	0.07	0.58	4
janos-us	0.093	0.108	0.63	7	0.43	2.65	1.55	30
janos-us-ca	0.141	0.138	0.83	8	0.80	5.76	2.76	243
newyork	0.018	0.034	0.14	1	0.07	0.05	0.28	1
nobel-eu	0.016	0.009	0.08	1	0.04	0.05	0.26	1
nobel-ger	0.011	0.020	0.04	1	0.04	0.08	0.24	3

SOCP models – Cplex (cont.)

nobel-us	0.015	0.083	0.10	1	0.04	0.04	0.19	1
norway	0.035	0.079	0.32	8	0.11	0.36	0.96	8
pdh	0.042	0.444	0.38	8	0.11	0.74	0.38	13
pioro40	0.019	0.039	0.27	1	0.10	0.14	0.57	6
polska	0.020	0.042	0.11	1	0.03	0.08	0.09	1
sun	0.165	0.587	0.89	13	0.65	7.68	2.36	257
ta2	0.020	0.015	0.13	1	0.12	0.08	0.89	4
garr 1999-01	0.022	0.017	0.13	1	0.09	0.21	0.33	1
garr 1999-04	0.029	0.000	0.07	0	0.10	0.07	0.45	3
garr 1999-05	0.029	0.004	0.09	1	0.10	0.08	0.40	3
garr 2001-09	0.030	0.000	0.10	0	0.11	0.10	0.44	3
garr 2001-12	0.029	0.000	0.08	0	0.09	0.16	0.32	3
garr 2004-04	0.028	0.000	0.18	0	0.09	0.05	0.31	3
garr 2009-08	0.087	0.005	0.46	2	0.57	0.47	1.99	27
garr 2009-09	0.089	0.011	0.62	4	0.60	0.61	2.19	36
garr 2009-12	0.090	0.013	0.78	4	0.60	0.59	2.47	44
garr 2010-01	0.093	0.013	0.50	4	0.61	0.57	2.32	32
w1-100-04	1.854	3.176	43.14	85	8.88	164.49	43.87	2585
w1-200-04	24.231	25.366	413.95	4075	231.09	2714.68	9088.54	127429

SOCP models – GUROBI

	GUROBI P				GUROBI bM			
	avg		max		avg		max	
	t	n	t	n	t	n	t	n
abilene	0.011	0.0	0.03	0	0.032	0.1	0.06	3
atlanta	0.012	0.5	0.03	8	0.044	1.6	0.08	15
cost266	0.012	0.4	0.05	11	0.099	0.8	0.30	27
dfn-bwin	0.007	0.0	0.01	0	0.068	0.0	0.08	0
dfn-gwin	0.017	0.0	0.04	0	0.104	0.1	0.31	4
di-yuan	0.028	2.0	0.21	46	0.116	4.9	0.46	74
france	0.011	0.3	0.03	6	0.079	1.2	0.18	17
geant	0.011	0.7	0.04	11	0.062	1.2	0.17	22
germany50	0.016	1.1	0.26	34	0.166	2.5	0.93	52
giul39	0.424	67.6	6.69	1308	1.795	138.5	30.02	2212
india35	0.014	0.4	0.12	14	0.132	1.8	0.34	29
janos-us	0.150	21.2	2.14	767	0.717	85.4	16.54	1168
janos-us-ca	0.285	47.1	7.87	916	1.741	158.4	25.93	1595
newyork	0.013	0.8	0.04	14	0.091	2.2	0.22	22
nobel-eu	0.013	0.2	0.09	9	0.080	0.4	0.25	31
nobel-ger	0.012	0.4	0.04	11	0.056	1.4	0.33	38

SOCP models – GUROBI (cont.)

nobel-us	0.012	0.8	0.05	11	0.047	0.9	0.15	11
norway	0.033	2.8	0.44	30	0.141	7.7	0.63	55
pdh	0.023	4.6	0.09	47	0.081	7.1	0.23	45
pioro40	0.015	0.6	0.09	13	0.160	2.6	0.57	44
polska	0.010	0.5	0.03	7	0.038	1.2	0.06	9
sun	0.189	39.6	0.76	282	0.961	126.9	5.68	583
ta2	0.018	0.6	0.12	27	0.214	1.9	1.52	33
garr 1999-01	0.034	0.5	0.09	9	0.096	6.6	0.38	17
garr 1999-04	0.016	1.9	0.11	26	0.115	2.7	0.55	35
garr 1999-05	0.018	2.0	0.08	25	0.139	3.5	0.79	36
garr 2001-09	0.020	2.0	0.09	19	0.156	4.0	0.97	29
garr 2001-12	0.015	0.0	0.04	0	0.116	0.1	0.31	17
garr 2004-04	0.021	3.0	0.06	14	0.128	3.5	0.57	27
garr 2009-08	0.070	7.6	0.72	124	0.776	18.8	5.39	164
garr 2009-09	0.071	7.6	0.59	202	0.918	21.8	4.85	212
garr 2009-12	0.071	7.6	0.55	123	0.920	22.7	6.21	352
garr 2010-01	0.073	7.6	0.68	114	0.916	22.8	5.76	339
w1-100-04	2.372	159.3	7.09	703	14.064	407.2	110.36	5339
w1-200-04	9.575	241.4	63.37	1395	134.145	637.0	2384.84	10943

Combinatorial properties

- A MINLP with strong network structure, how to exploit it?
- $\sigma = L = 0 \Longrightarrow r_{ij} = \rho x_{ij} \Longrightarrow$ Constrained Shortest Path (CSP) (this gives more than one idea, and proves $\mathcal{NP}$-hardness)
- Checking feasibility is easy: delay is a decreasing function of r_{ij} $\Longrightarrow$ simply push the rates to the maximum
- Modified arc costs $\bar{l}_{ij} = L/c_{ij} + L/w_{ij} + l_{ij} + n_i$
- For each $r \in C = \{ c_{ij} \ : \ (i,j) \in A \}$:
    - define reduced graph $G^r = (N, A^r)$ where $A^r = \{ (i,j) \in A \ : \ c_{ij} \geq r \}$
    - solve s-d shortest path P on G^r w.r.t. $\bar{l}$
    - if $\bar{l}(P) \leq \delta - \sigma / r$, then P feasible
    - if no feasible P found, then problem unfeasible (for fixed P, both LHS and RHS of (1) increase with r)
- Keep f-best solution found: (ERA-I heuristic)

Equal Rate Allocation

- Equal Rate Allocation: $r_{ij} = r\ (\geq \rho)$ for all $(i, j) \in P$ ($\Longrightarrow r_{min} = r$)
- It is easy to solve EC-ERA-SFSP-DCR for fixed r ($f_{ij} = 1$)
  - run Bellman-Ford on G^r with costs $l^r_{ij} = L/r + L/w_{ij} + l_{ij} + n_i$
  - each time d extracted from Q, check if delay $\leq \delta - \sigma/r$
- Repeating the above for all $r \in C$ does not solve EC-ERA-SFSP-DCR counterexample: returned path P with delay constraint not tight
- Obvious solution: for each P reduce r until constraint tight $\Longrightarrow$ keep feasibility, improve objective function $\Longrightarrow$ optimal
- Solves EC-ERA-SFSP-DCR, ERA-H heuristic for EC-SFSP-DCR
- Extensions:
  - Non-equal but integer f_{ij}: use dynamic programming instead of Bellman-Ford (pseudo-poly)
  - Non-equal continuous f_{ij}: classical approximation algorithm with cost rounding using the above pseudo-poly

ERA-Based Heuristics: Experiments

				ERA-I		ERA-H		
instance	n	m	k	avg	max	avg	max	inf
abilene	12	15	31	0.52	0.92	0.000	0.000	0.06
atlanta	15	22	45	0.57	0.88	0.000	0.000	0.07
cost266	37	57	120	0.48	0.95	0.000	0.000	0.17
dfn-bwin	10	45	45	0.03	0.06	0.000	0.000	0.00
dfn-gwin	11	47	53	0.16	0.86	0.000	0.000	0.02
di-yuan	11	42	58	0.48	0.90	0.000	0.000	0.12
france	25	45	66	0.44	0.90	0.000	0.000	0.02
geant	22	36	63	0.46	0.89	0.000	0.001	0.06
germany50	50	88	276	0.50	0.90	0.000	0.001	0.21
giul39	39	172	1482	0.67	0.97	0.011	0.570	0.10
india35	35	80	195	0.53	0.93	0.000	0.000	0.11
janos-us	26	84	650	0.71	0.95	0.004	0.275	0.18
janos-us-ca	39	122	1482	0.68	0.95	0.010	0.289	0.23
newyork	16	49	89	0.50	0.90	0.000	0.000	0.03
nobel-eu	28	41	106	0.55	0.93	0.000	0.000	0.23
nobel-ger	17	26	51	0.49	0.93	0.000	0.000	0.10

- inf = fraction of feasible wrongly declared unfeasible

ERA-Based Heuristics: Experiments (cont.)

nobel-us	14	21	24	0.35	0.90	0.000	0.001	0.00
norway	27	51	341	0.71	0.94	0.000	0.000	0.12
pdh	11	34	54	0.64	0.90	0.000	0.001	0.04
pioro40	40	89	204	0.40	0.89	0.000	0.000	0.25
polska	12	18	24	0.59	0.90	0.000	0.000	0.00
sun	27	102	702	0.76	0.95	0.008	0.431	0.06
ta2	65	108	388	0.45	0.92	0.000	0.000	0.31
garr 1999-01	16	36	240	0.65	0.88	0.000	0.001	0.02
garr 1999-04	23	50	506	0.57	0.94	0.000	0.001	0.75
garr 1999-05	23	50	506	0.55	0.94	0.000	0.000	0.75
garr 2001-09	22	48	462	0.60	0.94	0.000	0.000	0.74
garr 2001-12	24	52	552	0.59	0.94	0.000	0.000	0.75
garr 2004-04	22	48	462	0.56	0.94	0.000	0.000	0.75
garr 2009-08	54	136	2862	0.65	0.94	0.001	0.386	0.85
garr 2009-09	55	138	2970	0.67	0.94	0.000	0.000	0.85
garr 2009-12	54	136	2862	0.67	0.94	0.001	0.240	0.85
garr 2010-01	54	136	2862	0.67	0.94	0.001	0.241	0.85
w1-100-04	100	414	664	0.77	0.95	0.015	0.739	0.07
w1-200-04	200	1550	1528	0.71	0.96	0.015	0.814	0.05

3-Pronged Approach

- MI-SOCP approach accurate but slow
  ERA-* approaches fast but inaccurate
- Best of both worlds: 3-pronged approach
  1. run ERA-I, if instance unfeasible terminate
  2. otherwise run ERA-H: if a solution found, report it and terminate
  3. if all else fails, then run model P and report its solution
- So crude, does it really work?

3-Pronged Approach: Experiments

Cplex				GUROBI								
SOCP		3P		SOCP		3P		Gaps		ERA-H		
avg	max	avg	max	avg	max	avg	max	avg	max	avg	max	inf
0.009	0.02	0.001	0.01	0.009	0.02	0.001	0.01	0.00	0.00		0.00	0.06
0.016	0.16	0.001	0.02	0.010	0.03	0.001	0.02	0.00	0.00		0.00	0.07
0.013	0.05	0.002	0.03	0.012	0.04	0.003	0.04	0.00	0.00		0.00	0.17
0.011	0.02	0.000	0.00	0.007	0.01	0.000	0.01	0.00	0.00		0.00	0.00
0.019	0.09	0.000	0.01	0.015	0.04	0.000	0.01	0.00	0.00		0.00	0.02
0.050	0.35	0.017	0.35	0.028	0.22	0.012	0.23	0.00	0.00		0.00	0.12
0.015	0.04	0.000	0.01	0.010	0.03	0.000	0.01	0.00	0.00		0.00	0.02
0.013	0.05	0.001	0.01	0.010	0.04	0.001	0.03	0.00	0.00		0.00	0.06
0.021	0.09	0.005	0.08	0.017	0.24	0.007	0.27	0.00	0.00	7e-5	0.01	0.21
0.254	1.01	0.019	0.66	0.449	7.57	0.087	6.52	0.01	0.57	3e-4	0.01	0.10
0.019	0.25	0.002	0.04	0.016	0.11	0.002	0.07	0.00	0.00		0.00	0.11
0.091	0.62	0.013	0.33	0.153	2.25	0.051	2.19	0.00	0.28	1e-4	0.01	0.18
0.144	0.84	0.026	0.49	0.298	9.59	0.118	7.70	0.01	0.29	2e-4	0.01	0.23
0.017	0.13	0.000	0.02	0.015	0.04	0.001	0.02	0.00	0.00		0.00	0.03
0.014	0.05	0.004	0.05	0.016	0.09	0.005	0.09	0.00	0.00		0.00	0.23
0.010	0.03	0.002	0.03	0.015	0.04	0.002	0.04	0.00	0.00		0.00	0.10

3-Pronged Approach: Experiments (cont.)

0.013	0.09	0.000	0.00	0.014	0.05	0.000	0.00	0.00	0.00		0.00	0.00
0.032	0.30	0.005	0.25	0.035	0.32	0.005	0.13	0.00	0.00	6e-5	0.01	0.12
0.034	0.30	0.001	0.02	0.026	0.10	0.002	0.10	0.00	0.00		0.00	0.04
0.019	0.27	0.007	0.25	0.018	0.09	0.007	0.09	0.00	0.00	5e-5	0.01	0.25
0.016	0.09	0.000	0.00	0.014	0.03	0.000	0.00	0.00	0.00		0.00	0.00
0.154	0.89	0.006	0.36	0.188	0.87	0.009	0.40	0.01	0.43	2e-4	0.01	0.06
0.019	0.12	0.008	0.05	0.020	0.13	0.009	0.13	0.00	0.00	8e-5	0.01	0.31
0.025	0.12	0.001	0.03	0.035	0.10	0.001	0.03	0.00	0.00	4e-5	0.01	0.02
0.030	0.08	0.022	0.06	0.017	0.12	0.016	0.10	0.00	0.00	4e-5	0.01	0.75
0.028	0.08	0.021	0.06	0.018	0.08	0.016	0.08	0.00	0.00	6e-5	0.01	0.75
0.026	0.09	0.021	0.08	0.022	0.09	0.018	0.09	0.00	0.00	4e-5	0.01	0.74
0.027	0.07	0.022	0.07	0.016	0.04	0.012	0.04	0.00	0.00	4e-5	0.01	0.75
0.026	0.17	0.020	0.05	0.022	0.06	0.019	0.06	0.00	0.00	4e-5	0.01	0.75
0.084	0.44	0.075	0.44	0.069	0.70	0.065	0.71	0.00	0.39	2e-4	0.01	0.85
0.086	0.62	0.078	0.62	0.069	0.56	0.063	0.57	0.00	0.00	2e-4	0.01	0.85
0.088	0.75	0.078	0.73	0.071	0.52	0.061	0.50	0.00	0.24	2e-4	0.01	0.85
0.087	0.46	0.076	0.45	0.074	0.61	0.066	0.59	0.00	0.24	2e-4	0.01	0.85
1.906	46.7	0.034	1.84	2.354	8.35	0.150	3.54	0.01	0.74	2e-3	0.01	0.07
23.660	357.7	0.247	54.29	9.033	63.19	0.399	12.36	0.01	0.81	1e-2	0.02	0.05

Does it really matter in practice?

- Simulating the network behavior, large number of path computations
- Exponential interarrival (avg $= \lambda$), exponential duration (avg $=$ 1s)
- $\sigma = 3$ MTU and δ random in $[\, d_{min}\, , \, d_{min} + \beta(d_{max} - d_{min})\,]$
  $d_{min} =$ minimum feasible deadline, $d_{max} =$ delay constraint inactive
- Average of five independent replicas, and 95% confidence intervals
- Comparing all practical approaches known so far (2 new):
  1. ERA (equal rate allocation)
  2. SWPF-URA: shortest-widest-path + optimal (unequal) rate allocation
  3. WSPF-URA: widest-shortest-path + optimal (unequal) rate allocation
  4. SFSP-DCR: MI-SOCP model (perspective version)
  5. TPH: 3-pronged heuristic
- Same real-world topologies, realistic capacities

Simulation results: blocking probability

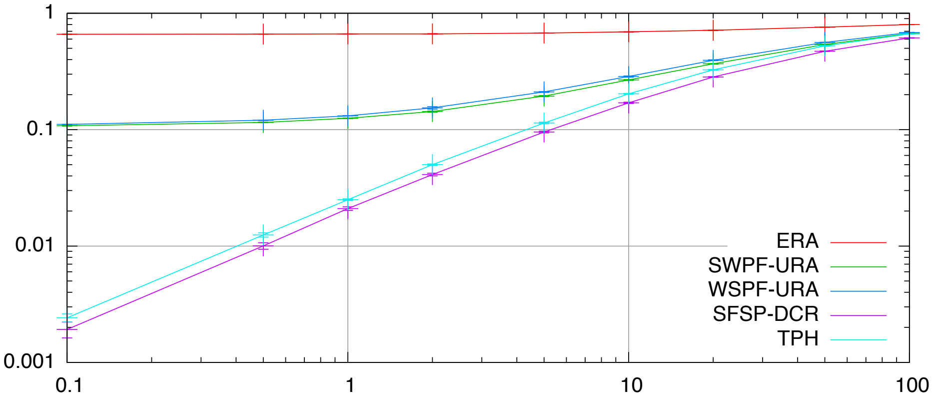

- ERA fails far too much (allocating the same rate a bad idea)
- both ERA and $*$-URA perform considerably worse than SFSP-DCR
- TPH performs quite close to the optimum
- Similar on all topologies, $\sigma \in \{1, 3, 10\}$MTU, $\beta \in \{0.2, 0.5, 1.0\}$

Why does ERA fail so often?

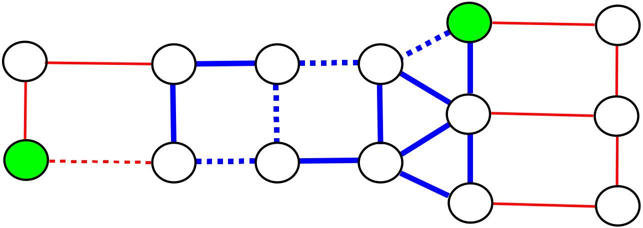

- Hub-and-spoke-like network with well-connected core (40/100 Gb) but weaker links to the periphery (1 Gb)
- Path from a core node to a peripheral one has to cross a weak link
- ERA has to allocate the same rate to all links $\Longrightarrow$ no more than the weak link's (residual) capacity $\Longrightarrow$ cannot meet the deadline
- The deadline can be met by reserving more capacity on core links

Simulation results: time

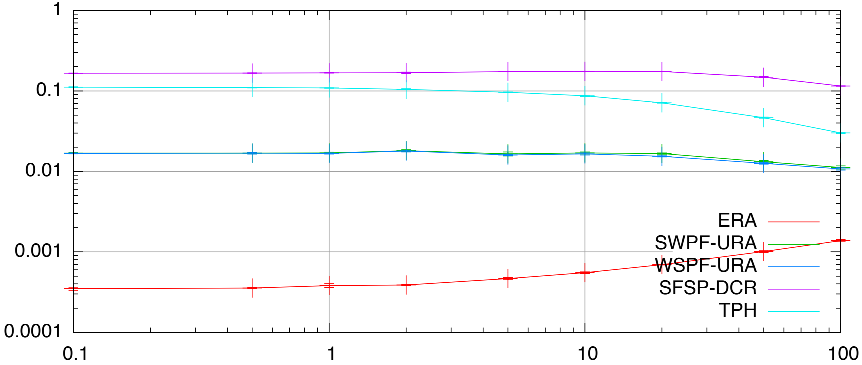

- SFSP-DCR slower but still affordable
- TPH much faster and almost as good
- “large” networks: $|N| = 70+$, $|A| = 230+$
- Path Computation Element makes this technically feasible

SOCP Model for WRP

- $\theta_{ij} = L/r_{ij} + |P(i,j)|L/w_{ij} \approx (2) \Longrightarrow$ (basically) same model
- But requires access control: not to make existing flows unfeasible
- Delay slack:
$$\bar{\delta}^k = \delta^k - \frac{\sigma^k}{r^k_{min}} - \sum_{(i,j)\in p^k} \left(\frac{L}{r^k_{ij}} + |P(i,j)|\frac{L}{w_{ij}} + l_{ij} + n_i \right)$$
- Access control constraint, one for each $k \in K$
$$\sum_{(i,j)\in p^k} \frac{L}{w_{ij}} x_{ij} \leq \bar{\delta}^k$$
- Can be used to preprocess away some arcs
- The coefficients are the same for each flow, can use path (+ RHS) dominance to detect redundant ones
- Still, possibly many constraints ($|K| \approx n^2$)

SOCP Model for FB

- $\theta_{ij} = L/r_{ij} + (|P(i,j)| + \phi(r_{ij}))L/w_{ij}$ ($\approx$ WRP) where

$$\phi(r) = (w_{ij} - r)/\min\{\, r \,,\, \bar{r}_{ij} \,\}$$

- Since $\bar{r}_{ij}$ is fixed, can be rewritten as

$$\phi(r) = \begin{cases} \phi_1(r) = w_{ij}/r - 1 & \text{if } 0 < r \leq \bar{r}_{ij} \\ \phi_2(r) = (w_{ij} - r)/\bar{r}_{ij} & \text{if } \bar{r}_{ij} \leq r \leq c_{ij} \; (\leq w_{ij}) \end{cases}$$

- Convex!: $\phi'_1(\bar{r}_{ij}) \leq \phi'_2(\bar{r}_{ij})$

- Can use the classical variable splitting reformulation

$$\theta_{ij} = v_{ij} + v'_{ij} + \frac{L}{w_{ij}} \left[|P(i,j)| x_{ij} - \frac{r_{ij}}{\bar{r}_{ij}} \right]$$

$$r_{ij} = r'_{ij} + r''_{ij} \quad , \quad \rho x_{ij} \leq r'_{ij} \leq \bar{r}_{ij} x_{ij} \quad , \quad 0 \leq r''_{ij} \leq (c_{ij} - \bar{r}_{ij}) x_{ij}$$

$$v_{ij} r_{ij} \geq L x_{ij}^2 \;\; , \;\; v_{ij} \geq 0 \;\; , \;\; v'_{ij} r'_{ij} \geq L x_{ij}^2 \;\; , \;\; v'_{ij} \geq 0$$

two conic constraints to represent the same L/r_{ij}: can we do better?

SOCP Model for FB (cont.d)

- Actually we can: $\phi_1(w_{ij}) = \phi_2(w_{ij}) \Longrightarrow \phi(r) = \max\{\, \phi_1(r) \,,\, \phi_2(r) \,\}$!
- Can use the "cutting planes" representation of ϕ:

$$\theta_{ij} = v_{ij} + v'_{ij} + \frac{L}{w_{ij}}(|P(i,j)| + 1)x_{ij} \;\;,\;\; v_{ij} r_{ij} \geq L x_{ij}^2 \;\;,\;\; v_{ij} \geq 0$$

$$v'_{ij} \geq v_{ij} - L/w_{ij} \;\;,\;\; v'_{ij} \geq (L/\bar{r}_{ij})x_{ij} - L r_{ij}/(w_{ij}\bar{r}_{ij})$$

  only one conic constraint, less variables (is this better?)

- Admission control constraint for FB:

$$\sum\nolimits_{(i,j)\in p^k} \frac{L}{w_{ij}}(x_{ij} + (w_{ij} - r^k_{ij})z_{ij}) \leq \bar{\delta}^k$$

$$s_{ij} \leq r_{ij} \;\;,\;\; s_{ij} \leq \bar{r}_{ij} \;\;,\;\; s_{ij} z_{ij} \geq x_{ij}^2 \;\;,\;\; z_{ij} \geq 0$$

  $+2|A|$ variables, $+|A|$ conic constraints but shared among flows

- Different coefficients (to share the z_{ij}), dominance more difficult
- Arc-based preprocessing still possible (using $r_{ij} = c_{ij}$)

Computational results for WRP and FB

Computational results for WRP and FB

Er . . . , not ready yet, sorry!

Extending the combinatorial approaches to WRP and FB

Extending the combinatorial approaches to WRP and FB

Er . . . , not ready yet, sorry!

Conclusions

- The world is indeed nonlinear, but surprisingly often nicely convex
- DCR: interesting generalization of classical "steady state" flows
- Relevant for applications, apparently good results
- MI-SOCP with substantial network structure = prototypical blend of nonlinear and combinatorial optimization
- MINLP techniques useful (Perspective Reformulation, SOCP, . . .)
- Combinatorial techniques useful (shortest paths, dynamic programming, approximation algorithms, . . .)
- Both are needed
- Still lots of work to do (WRP/FB, multi-flow, multi-path, network design, robust, . . .), problems look pretty hard
- Lots of fun. Join in! :-)