Traffic Engineering with Estimated Traffic Matrices

Matthew Roughan www.research.att.com/~roughan
Mikkel Thorup www.research.att.com/~mthorup
Yin Zhang www.research.att.com/~yzhang

AT&T Shannon Lab

Can we do IP route optimization?

- Feldmann et al., 2000
- Shaikh et al., 2002

- Zhang et al., 2003
- Zhang et al., 2003

- Fortz et al., 2002

A3: "Well, we don't know the topology, we don't know the traffic matrix, the routers don't automatically adapt the routes to the traffic, and we don't know how to optimize the routing configuration. But, other than that, we're all set!"

"A Northern New Jersey Research Lab"

Problem

- How well do all of these things work together?
  - If we do TE based on estimated TMs, how well do the results perform on the real TM?
- Question 1
  - Traffic matrices can be estimated from link data
  - How important are estimation errors?
  - Simple statistics don't tell the whole story!
- Question 2
  - Route optimization assumes good input data
  - How robust are different methods to input errors?

Methodology

- Need realistic test
  - Simulations can produce anything we want
  - Need realistic TMs and errors
    - Random errors quite different from systematic
  - Need realistic network
- Use data from AT&T's backbone
  - Topology, and 80% TM from Cisco Netflow
- Use existing techniques (as blackboxes)
  - TM Estimation
  - Route optimization (example of TE)
- Approach
  - apply optimizer on estimated TMs
  - test performance on actual TMs

Refresher: TM estimation

Have link traffic measurements (from SNMP)
Want to know demands from source to destination

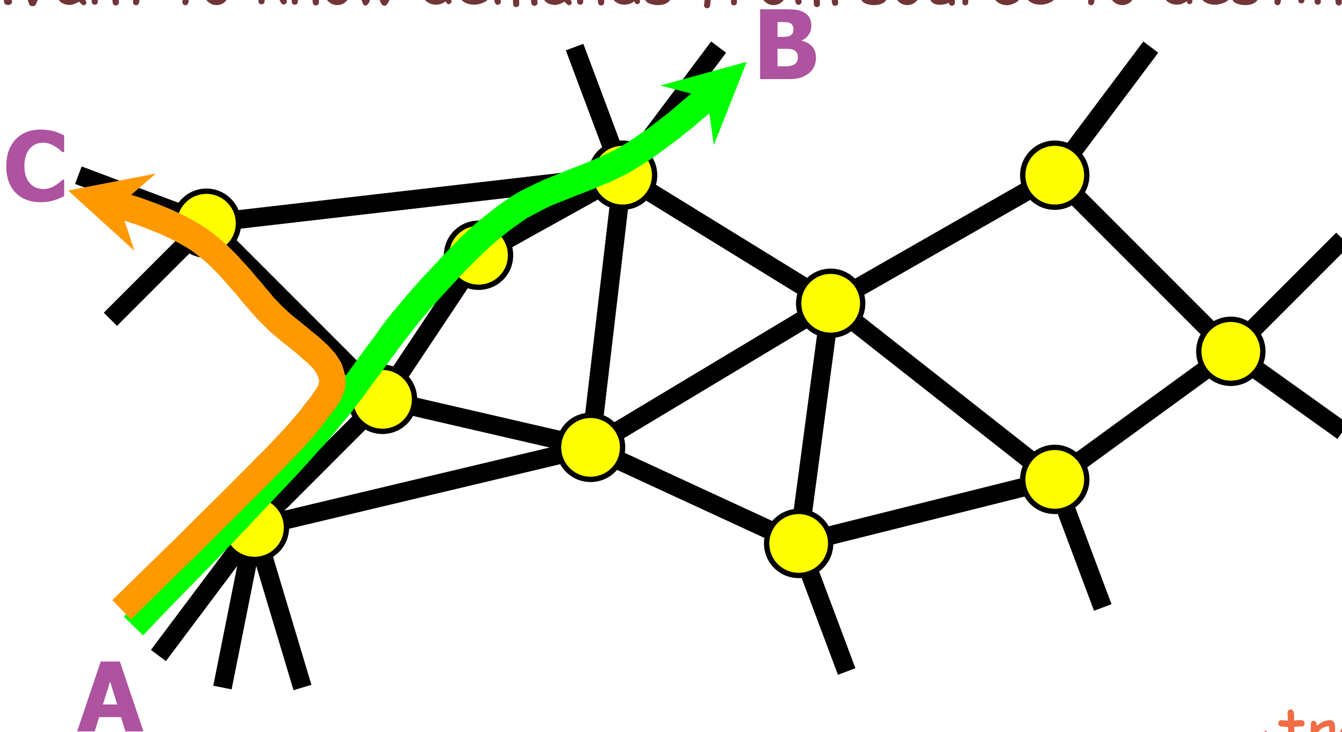

traffic matrix

$$x = At$$

link measurements

routing matrix

Methods of TM estimation

- Gravity Model
  - Demands are proportional
- Generalized Gravity Model
  - Take into account hot-potato routing asymmetry
- Tomo-gravity combines
  - Internal (tomographic) link constraints: $x=At$
  - Generalized gravity model
- Other methods
  - Not tested here (MLE, and Bayesian approaches)
  - Hard to implement at large scale
    - TE requires at least router-router TM

TM estimation results

Average traffic + |Error|

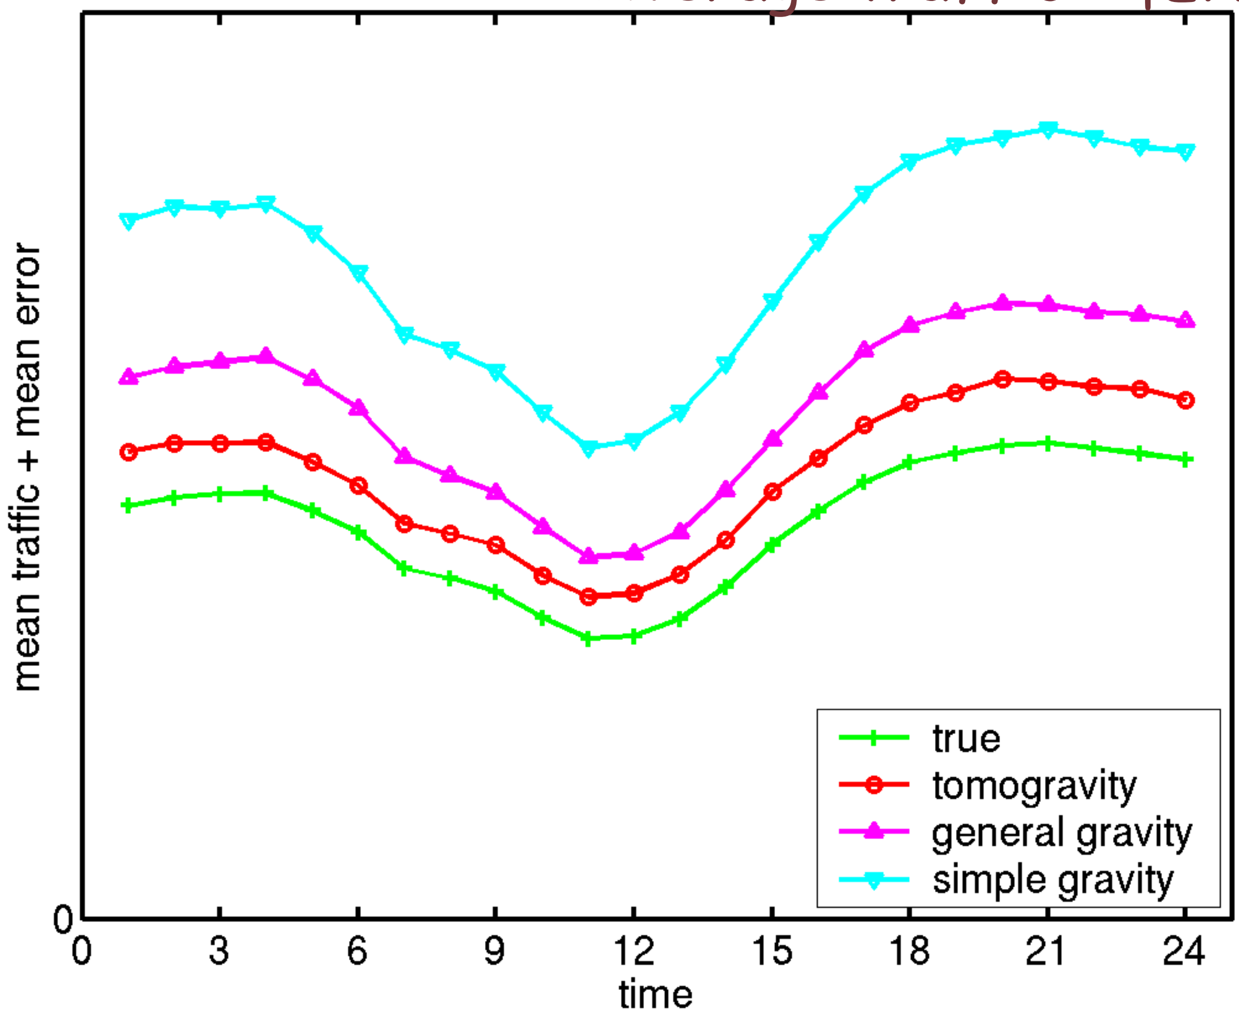

Route Optimization

- Route Optimization
  - Choosing route parameters that use the network most efficiently
  - Measure efficiency by maximum utilization
- Methods
  - Shortest path IGP weight optimization
    - OSPF/IS-IS
    - Choose weights
  - Multi-commodity flow optimization
    - Implementation using MPLS
    - Arbitrary splitting of traffic
    - Explicit set of routes for each origin/destination pair

Methodology: details

1. Start with real TM: t
   - ➔ measured using netflow
2. Simulate link measurements: $x = At$
3. Estimate TM: $\hat{t}(x)$
   - ➔ Use gravity/tomogravity methods
4. Compute optimal routing: $\hat{A}(\hat{t})$
   - ➔ Use MPLS/OSPF methods
5. Apply routing matrix Â on real TM $\hat{x} = \hat{A}t$
6. Compute

$$\text{max-utilization}(\hat{t};t) = \max_i \frac{\hat{x}_i}{c_i}$$

Results

OSPF + tomo-gravity good!

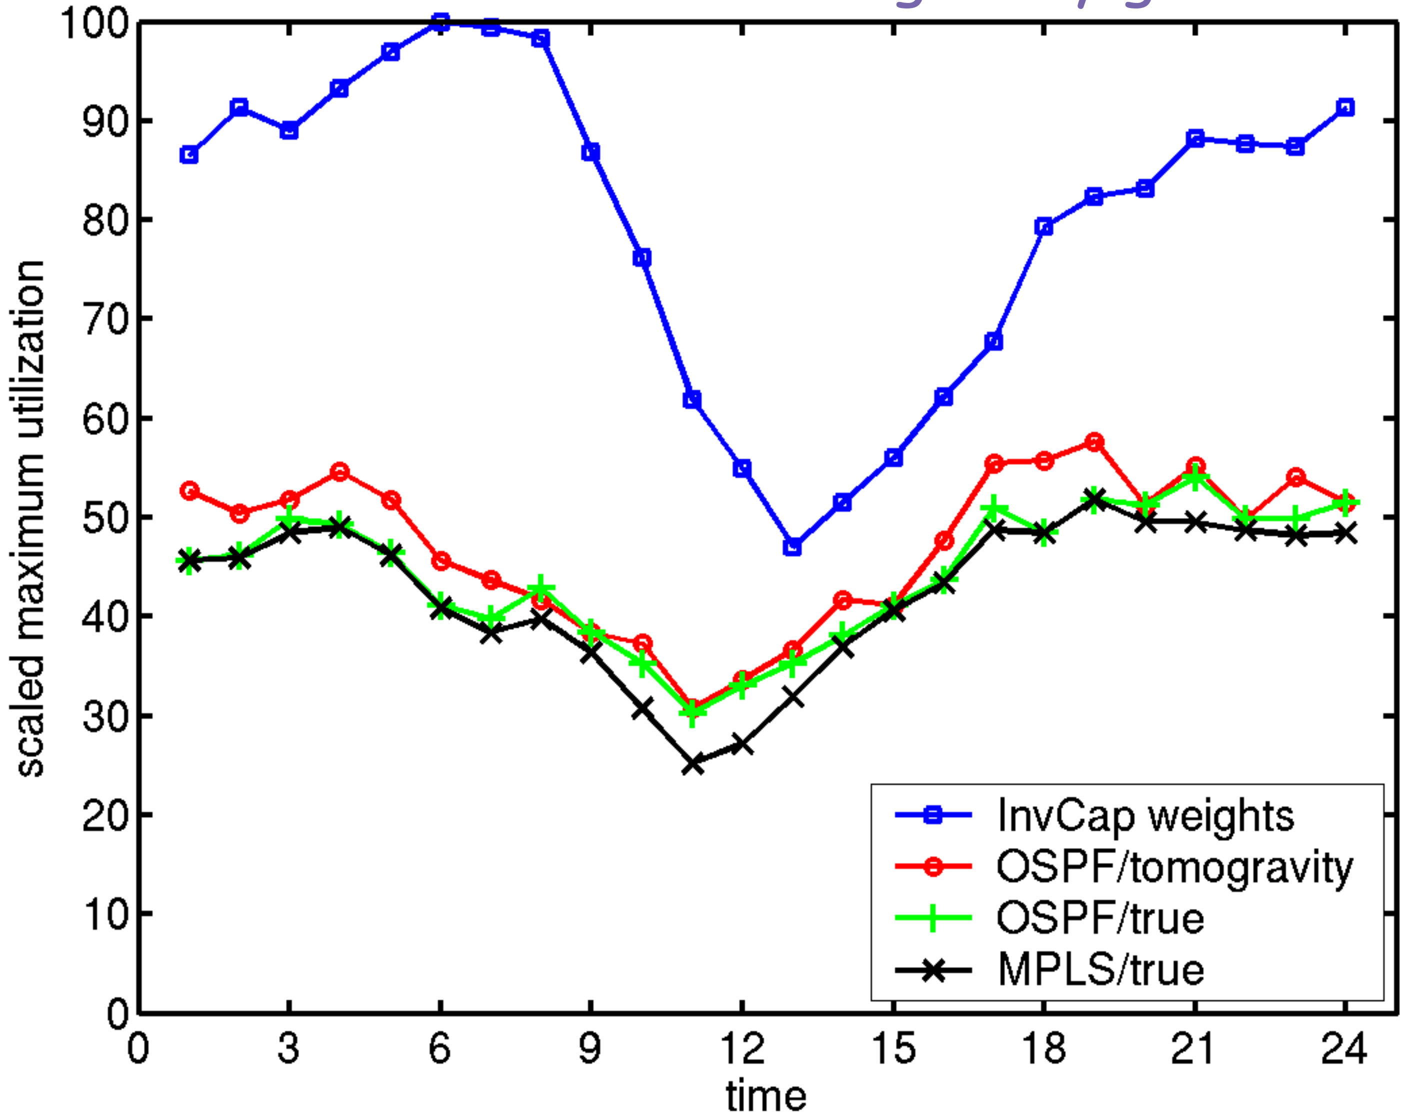

Results

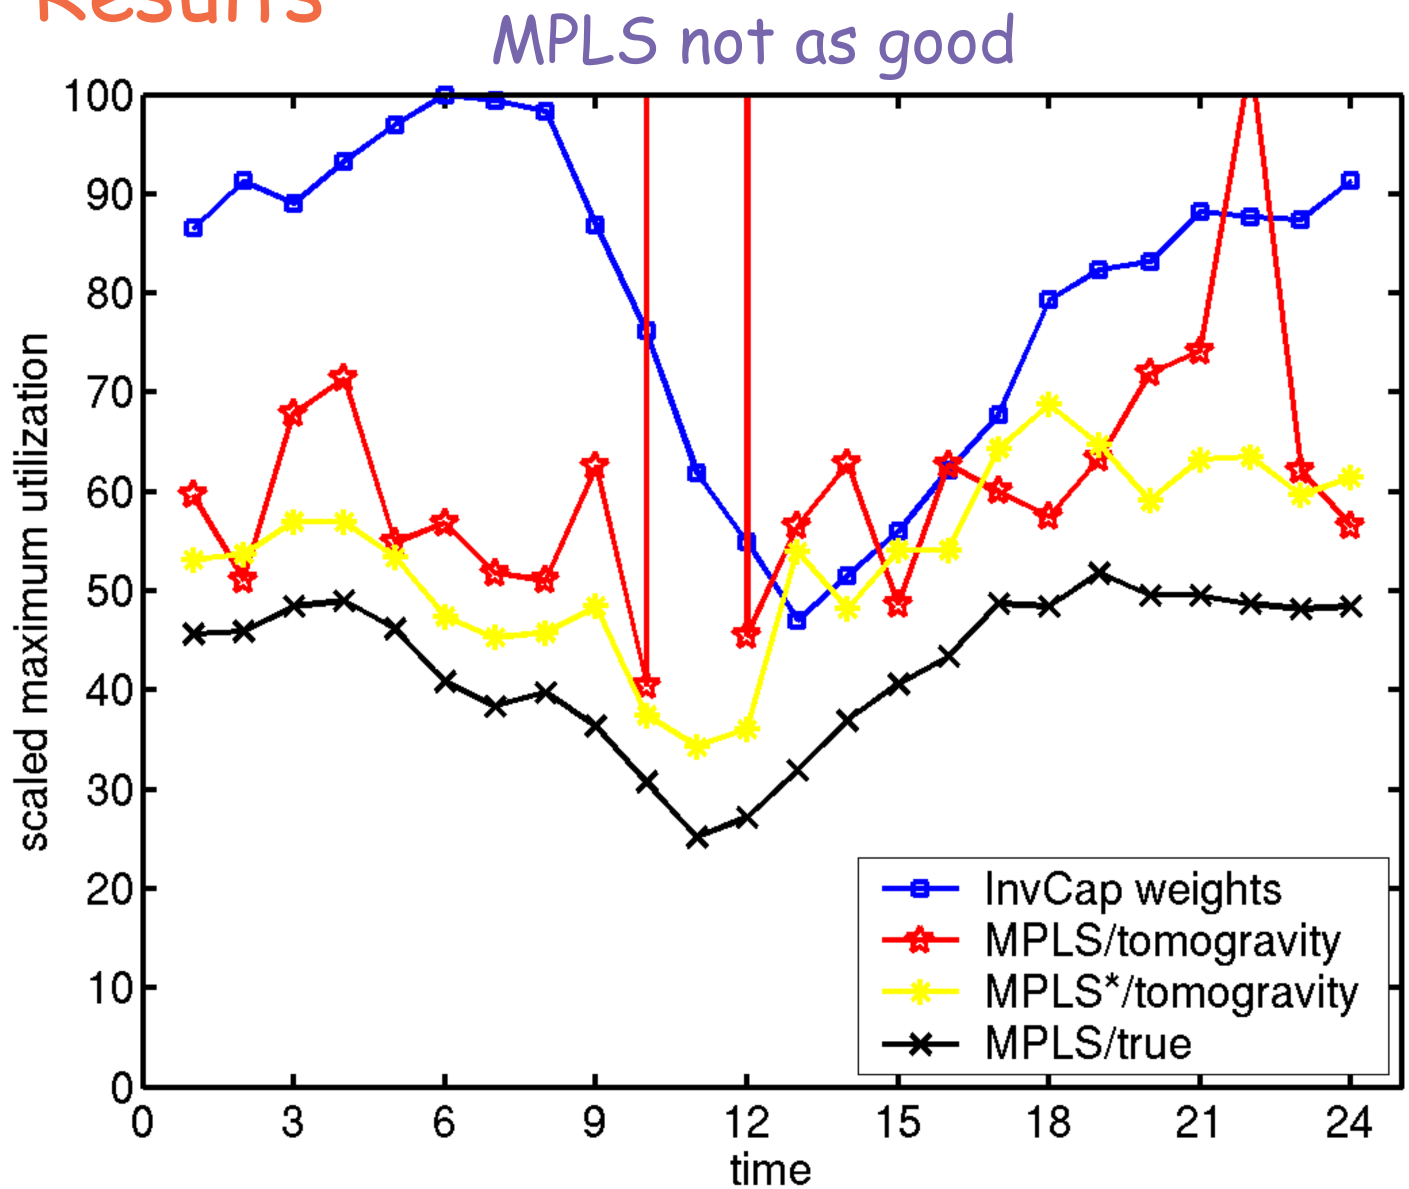
MPLS not as good
scaled maximum utilization
time
0
10
20
30
40
50
60
70
80
90
100
0
3
6
9
12
15
18
21
24
InvCap weights
MPLS/tomogravity
MPLS*/tomogravity
MPLS/true

Results

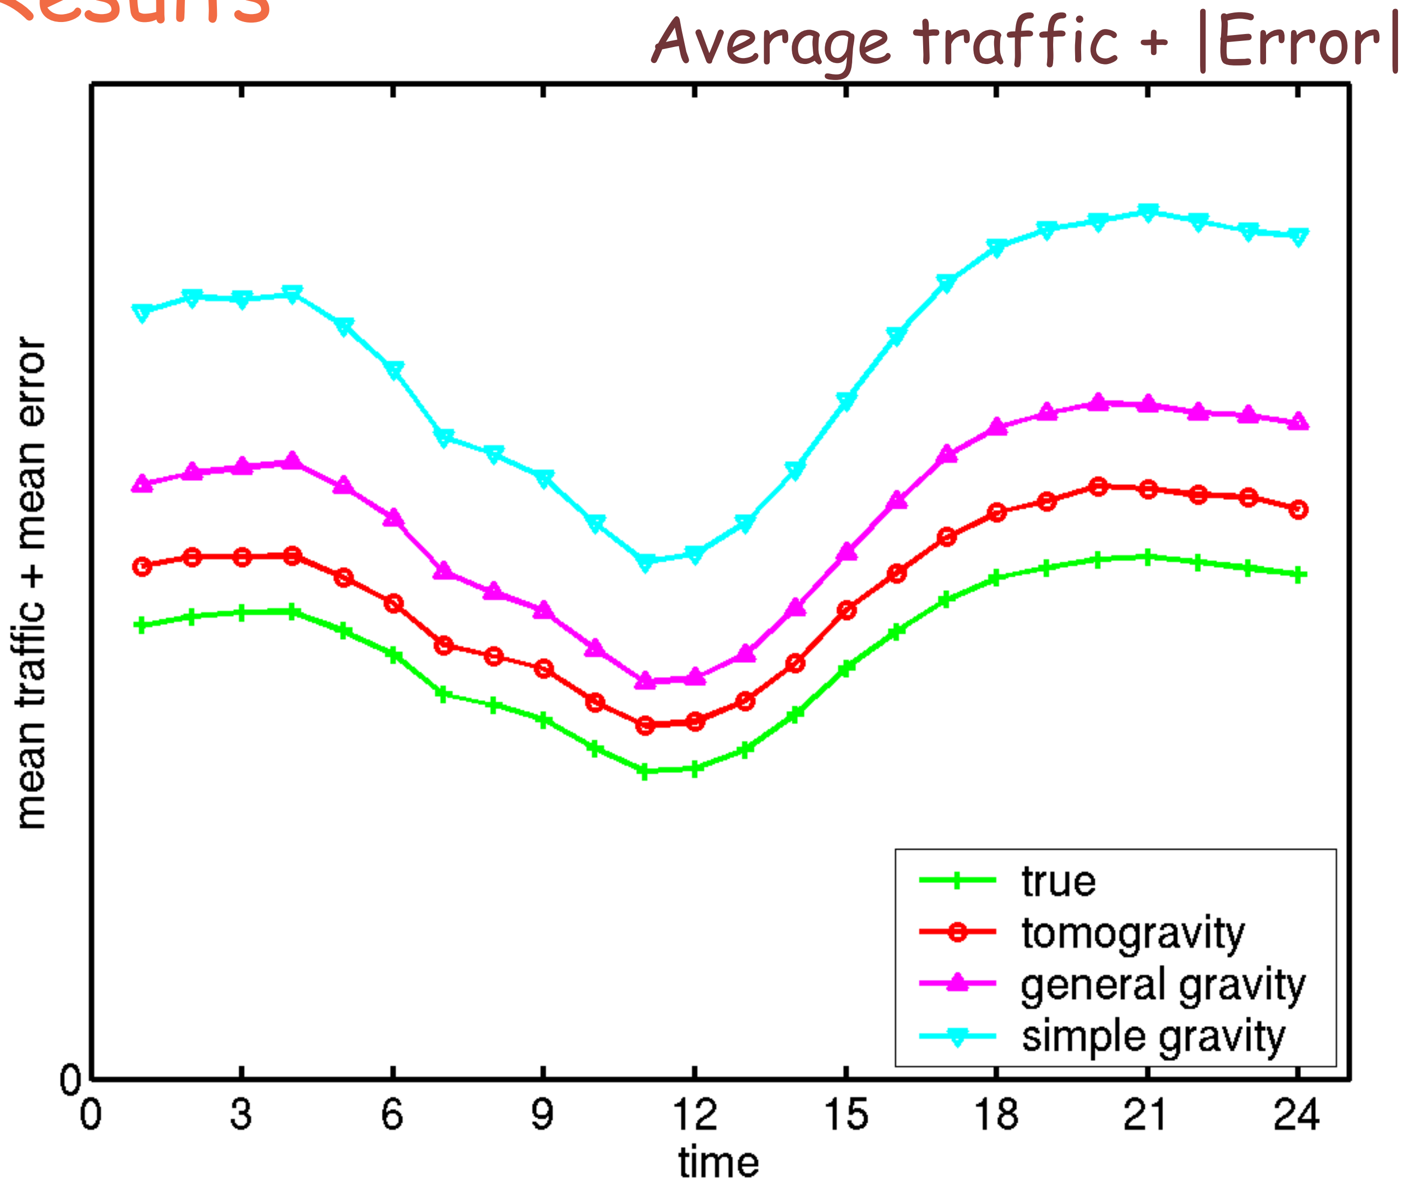
Average traffic + |Error|
mean traffic + mean error
time
0
0
3
6
9
12
15
18
21
24
true
tomogravity
general gravity
simple gravity

Results

Errors in TM: magnitude is not the key

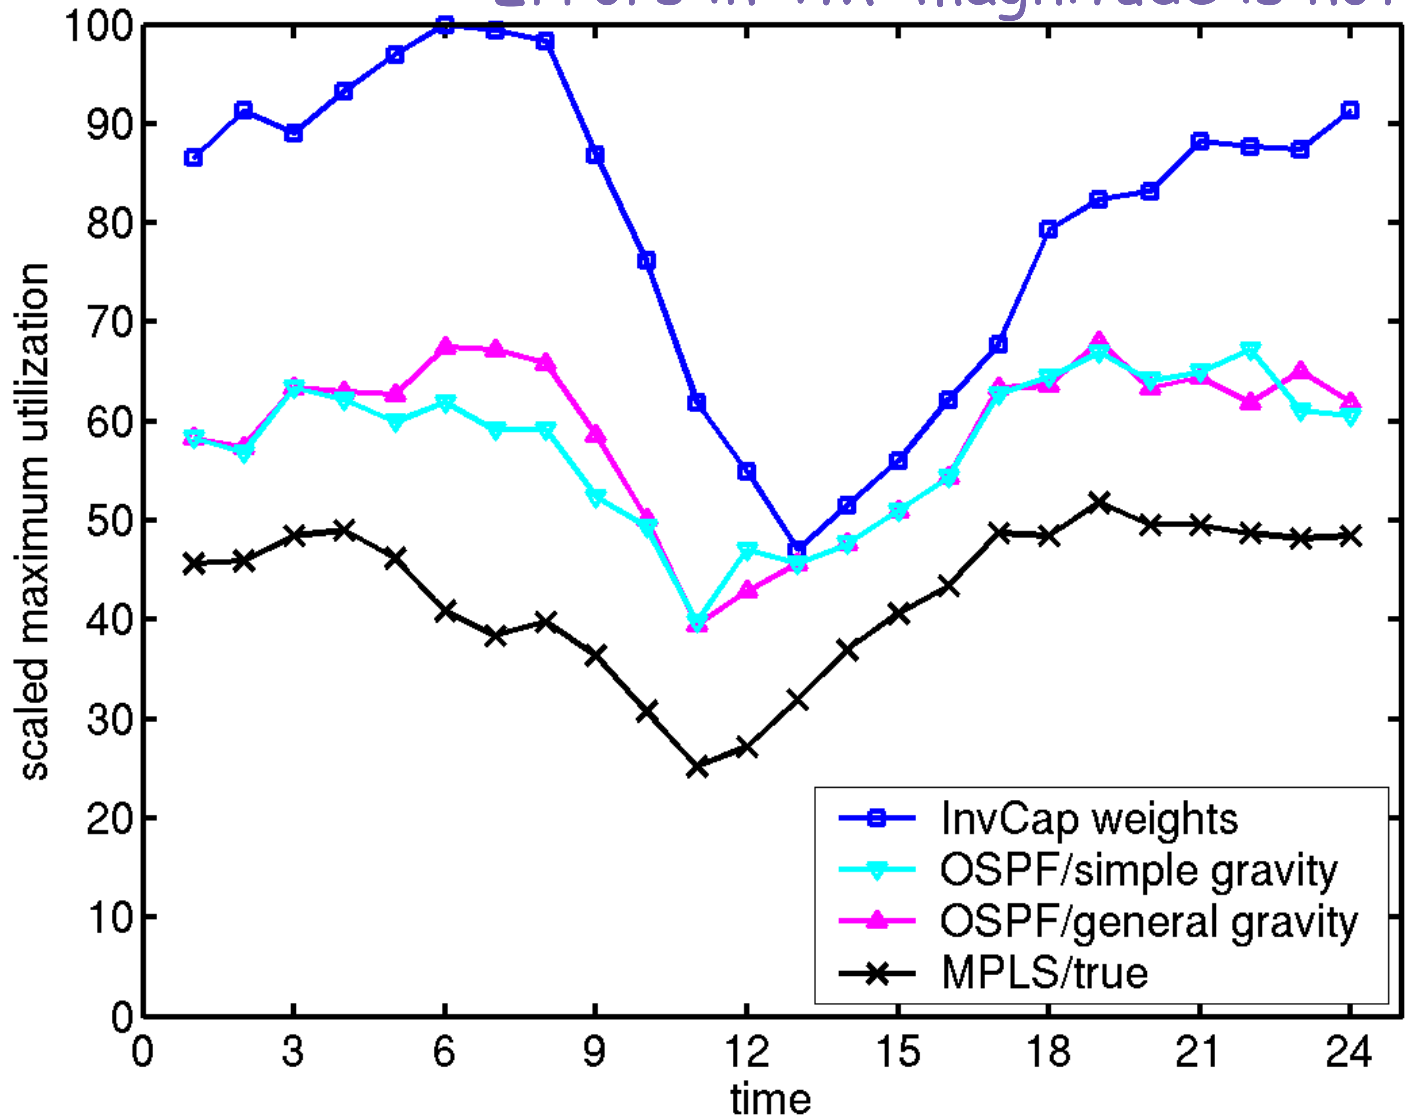

Other properties

* Other utility functions
* Global Optimization
  * Can optimize OSPF weights for 24 (1 hour) TMs
* Predictive mode
  * Works up to 7 days (at least)
* Fast convergence
  * Don't need as many iterations if speed is important
* Can design for limited no. of weight changes
  * Much of benefit from a few changes
* Can design for failure scenarios
  * Weights that work well for normal + failure modes

Conclusion

- Important to study TE and TM errors together
  - ➔ Simple statistics of errors don't indicate results
  - ➔ Best optimizer doesn't work best with input errors
  - ➔ Note: even measured TMs are used predictively
- TM Estimation and route optimization can work well together
  - ➔ IGP weight optimization
  - ➔ Robust
  - ➔ Close to optimum
  - ➔ Stable (predictive performance)

Acknowledgements

- Data collection
  - Fred True, Joel Gottlieb, Carsten Lund
- Tomogravity
  - Albert Greenberg and Nick Duffield
- OSPF Simulation
  - Carsten Lund, Nick Reingold

Additional Slides

TM estimation results

Relative mean error

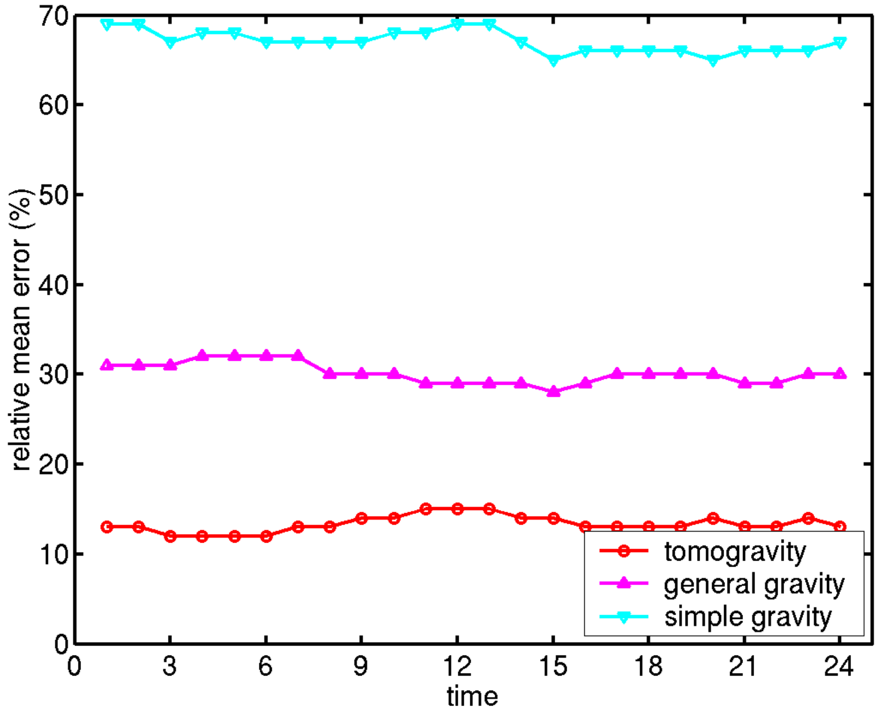

Average value + |Error|

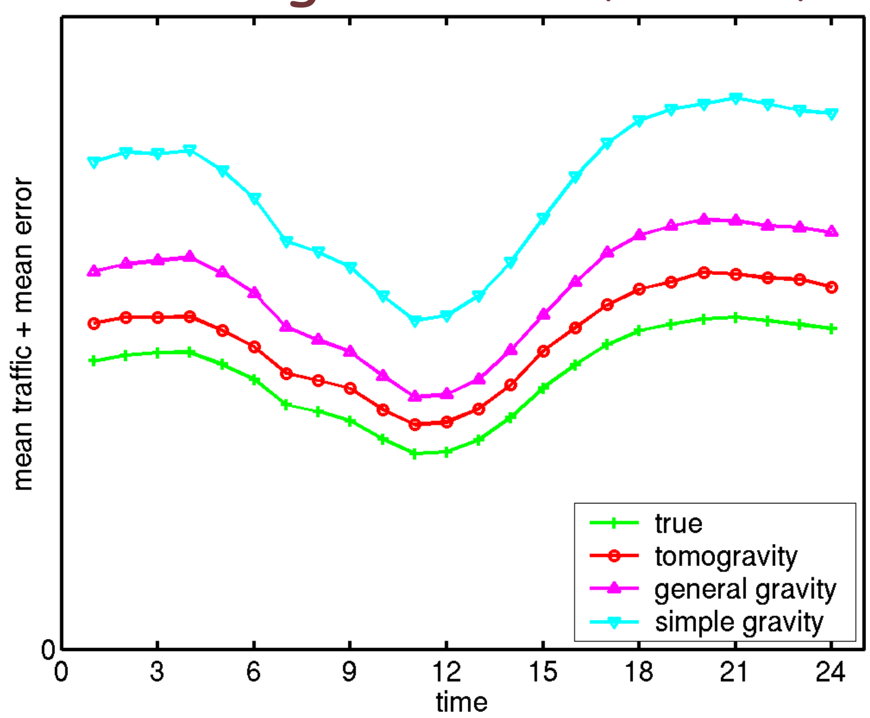

Global optimization