SNMP

Simple Network Measurements Please!

Matthew Roughan (+many others)

<roughan@research.att.com>

Outline

- Part I: SNMP traffic data
  - Simple Network Management Protocol
  - Why? How? What?
- Part II: Wavelets
  - What can you do?
  - Why not?
- Part III: Modeling
  - Putting time series and traffic modeling together
    - Traffic modeling deals with stationary processes (typically)
    - Time series gives us a way of getting a stationary process
    - But the analysis requires an understanding of the traffic model

Part I: SNMP Traffic Data

Data Availability – Traffic Data

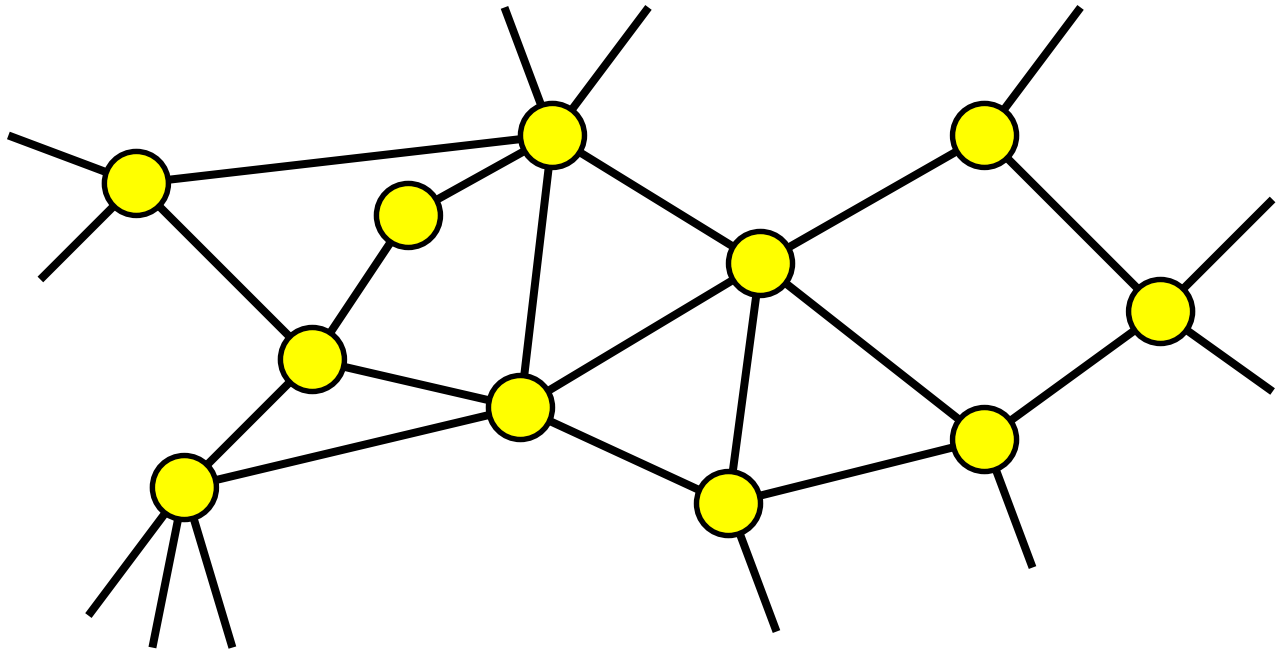

Data Availability – packet traces

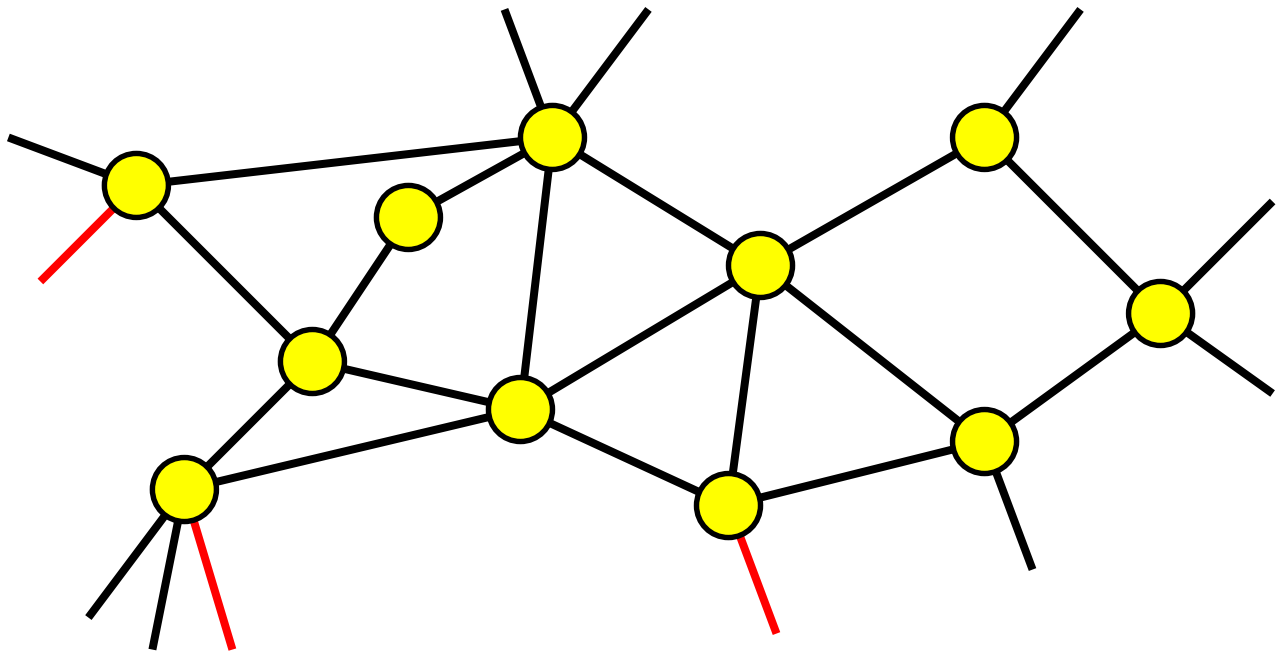

Packet traces limited availability

- special equipment needed (O&M expensive even if box is cheap)
- lower speed interfaces (only recently OC48 available, no OC192)
- huge amount of data generated

Data Availability – flow level data

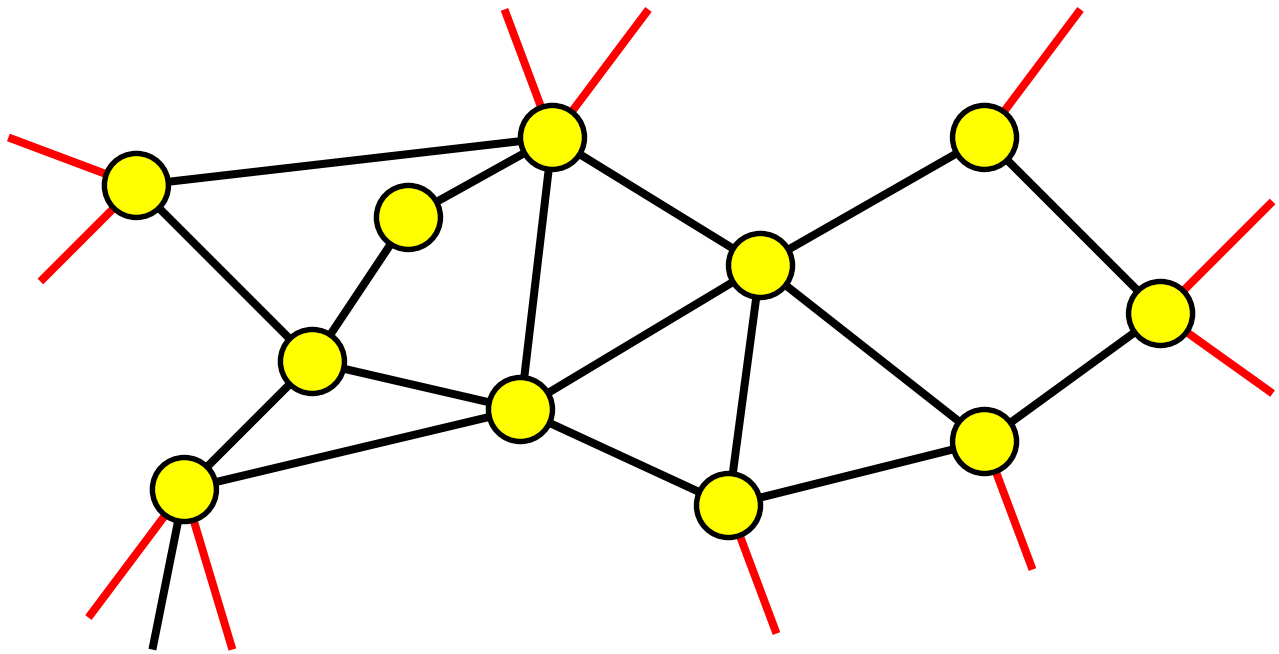

Flow level data not available everywhere

- historically poor vendor support (from some vendors)
- large volume of data (1:100 compared to traffic)
- feature interaction/performance impact

Data Availability – SNMP

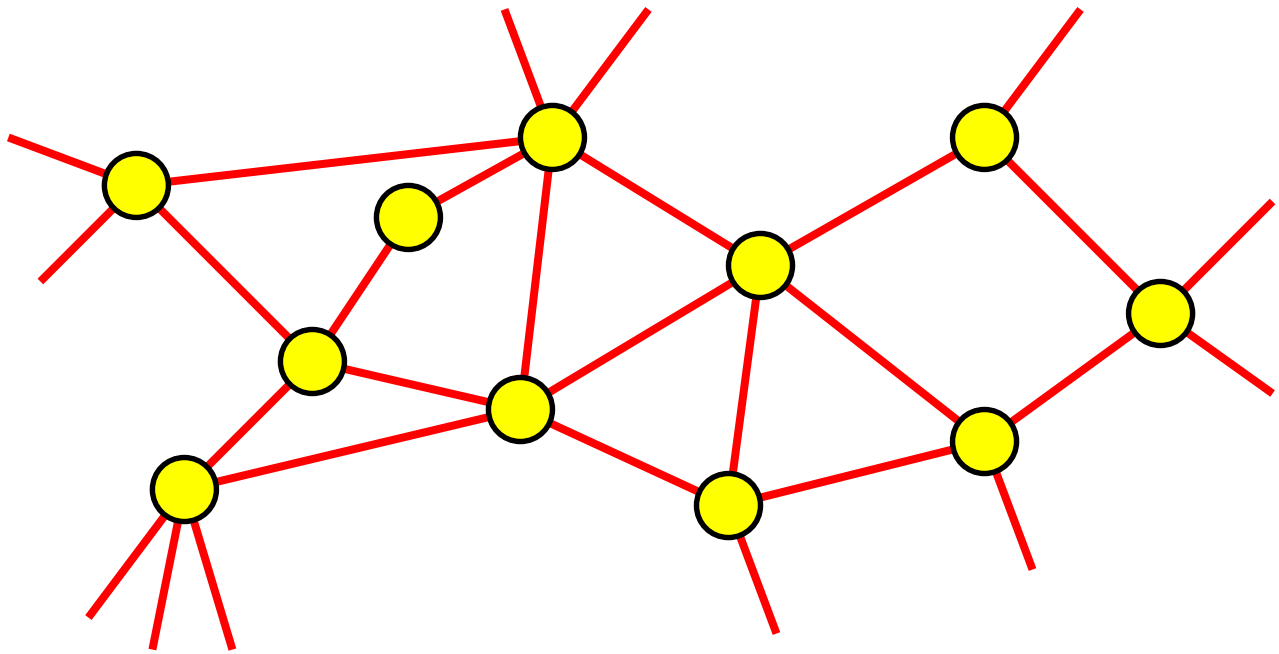

SNMP traffic data

- MIB II (including IfInOctets/IfOutOctets) is available almost everywhere
- manageable volume of data
- no significant impact on router performance

SNMP

- Advantages (MIB-II: IfInOctets/IfOutOctets)
  - Simple, Easy, available anywhere that supports SNMP
  - Relatively low volume
  - It is used by operations already (lots of historical data)
- Disadvantages
  - Data quality
    - Ambiguous
    - Missing data
    - Irregular sampling
  - Octets counters only tell you link utilizations
    - Hard to get a traffic matrix
    - Can't tell what type of traffic
    - Can't easily detect DoS, or other unusual events
  - Coarse time scale (>1 minute typically)
    - Lack of well tested relationship between coarse time-scale averages and performance (hence active perf. measurement)

SNMP traffic data

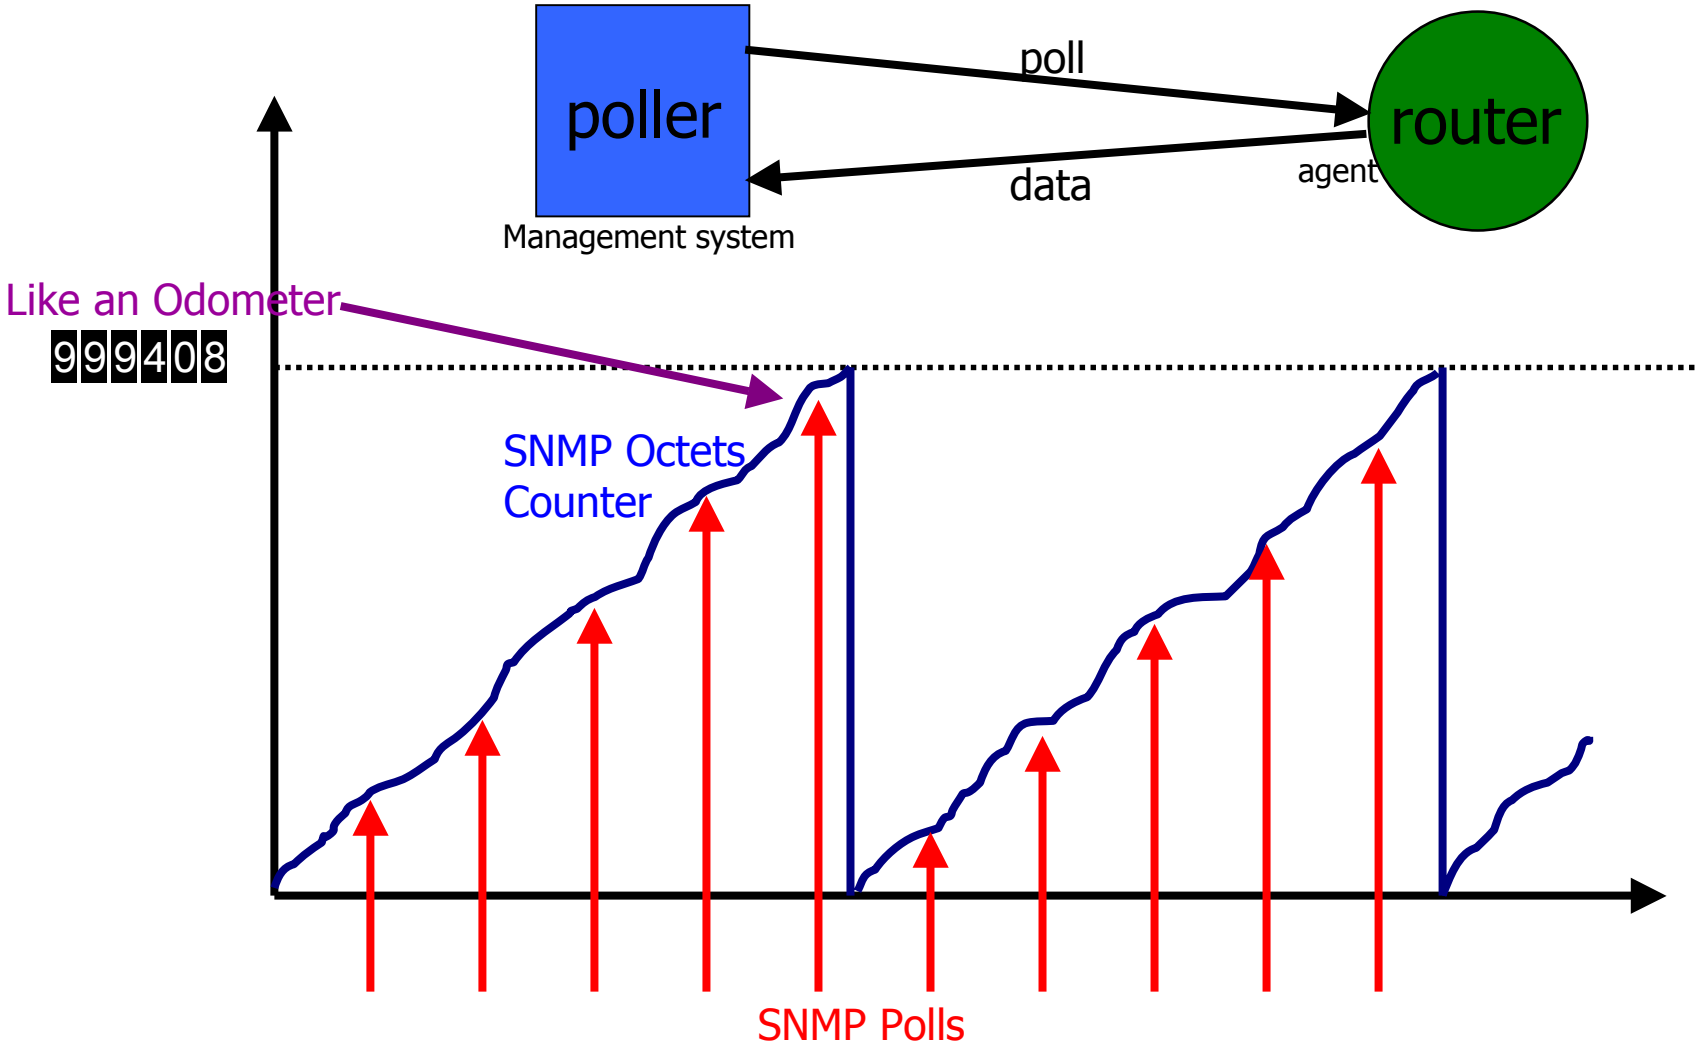

Irregularly sampled data

- Why?
  - Missing data (transport over UDP, often in-band)
  - Delays in polling (jitter)
  - Poller sync
    - Multiple pollers
    - Staggered polls

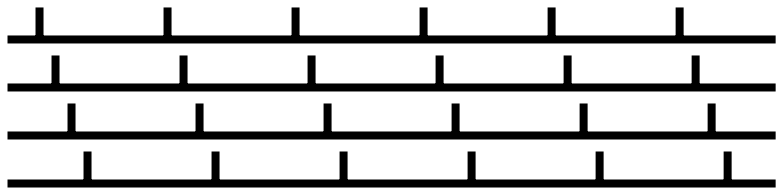

- Why care?
  - Time series analysis
  - Comparisons between links
    - Did traffic shed from link A go to link B
    - Calculation of traffic matrices
  - Totals (e.g. total traffic to Peer X)
  - Correlation to other data sources
    - Did event BGP route change at time T effects links A,B,C,…

Applications

- Capacity planning
  - Network at the moment is "hand-crafted"
  - Want to automate processes
  - Provisioning for failure scenarios requires adding loads
- Traffic engineering
  - Even if done by hand, you need to see results
  - BGP
- Event detection
  - Operations are "fire-fighters"
  - Don't care about events if they go away
  - Don't see patterns
- Business cases
  - Help sales and marketing make cases

Part II: Wavelet Analysis

- Multi-scale
- Multi-resolution

Discrete Wavelet Transform

- Replace sinusoidal basis functions of FFT with wavelet basis functions
- Implementation in pyramidal filter banks

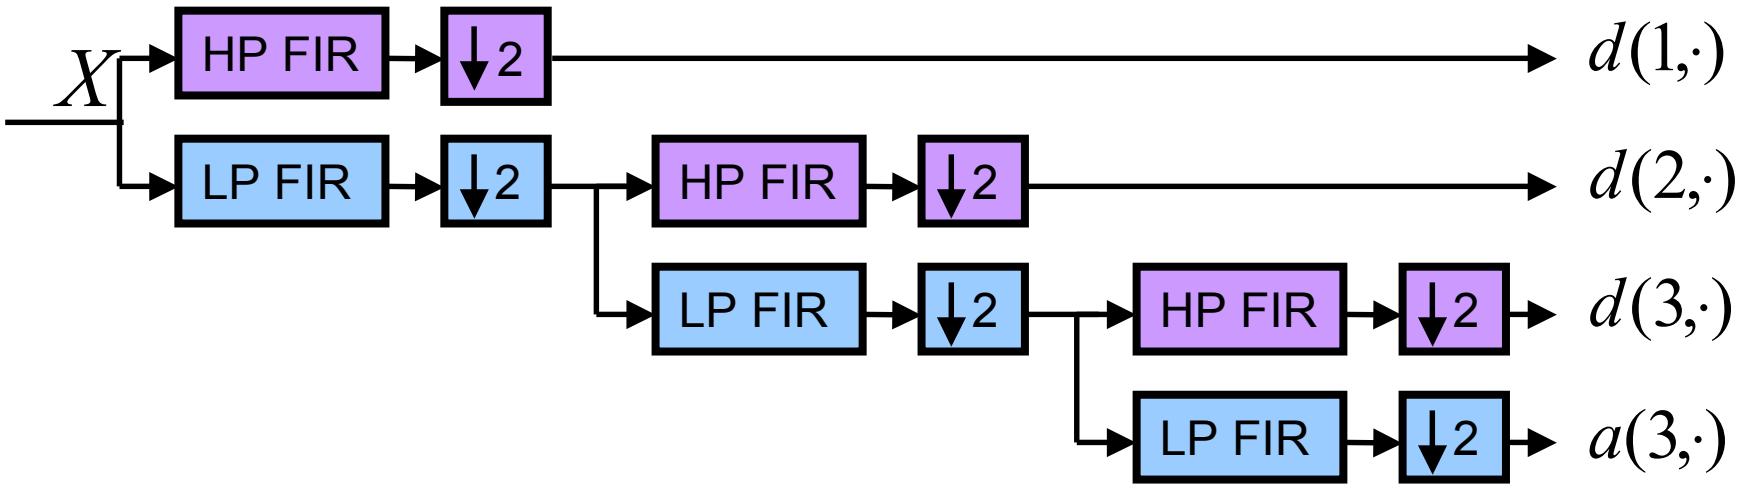

Dyadic grid

- no redundancy, no loss of information
- Each frequency/scale examined at a resolution matched to its scale

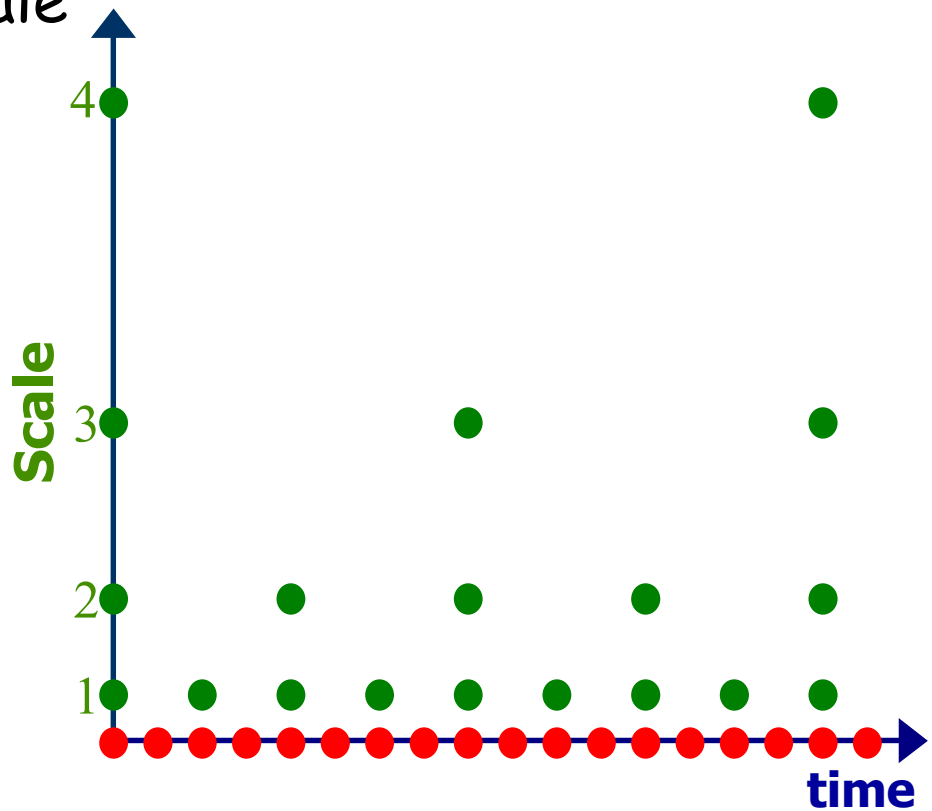

Dyadic grid: smoothing

- Zero the fine scale details and reconstruct

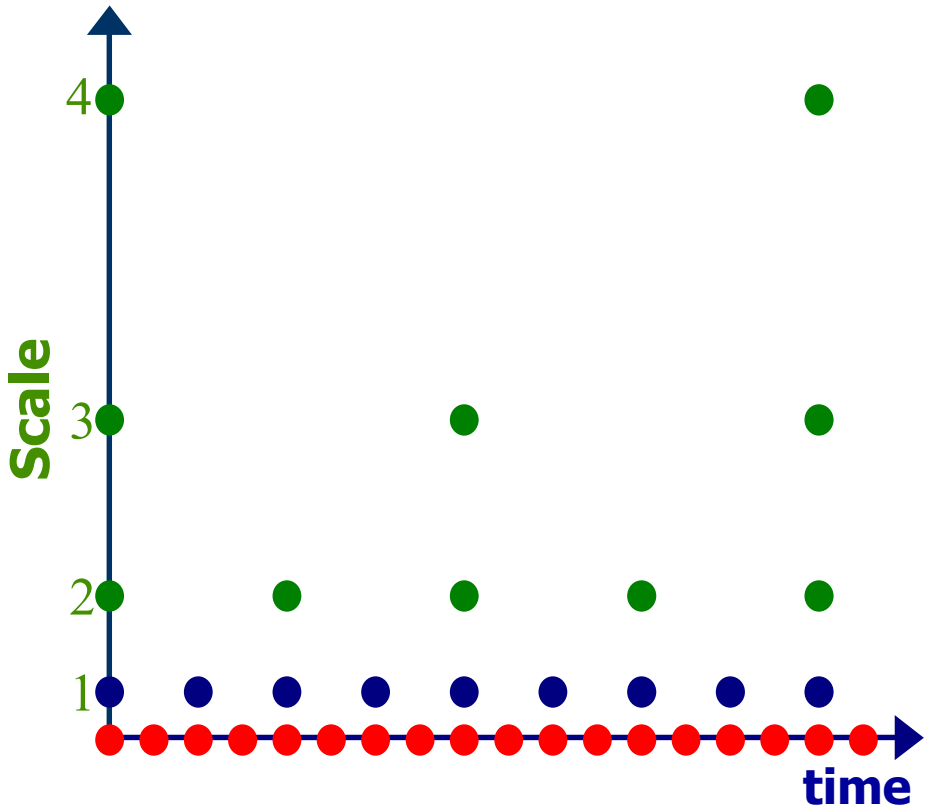

Dyadic grid: compression

- Keep the coefficients above some threshold

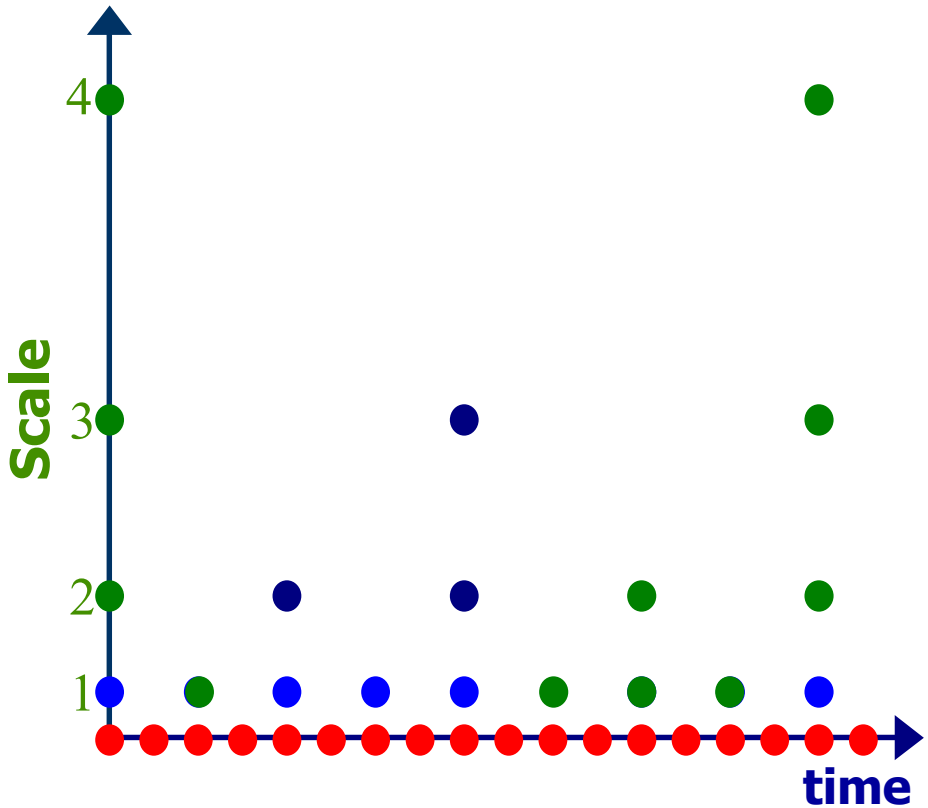

What can you do with wavelets

- Compression
- Smoothing/interpolation
- Anomaly detection/identification
  - DoS
  - Flash crowds
- Multiple dimensional analysis of data
- LRD/self-similarity analysis

Example: compression

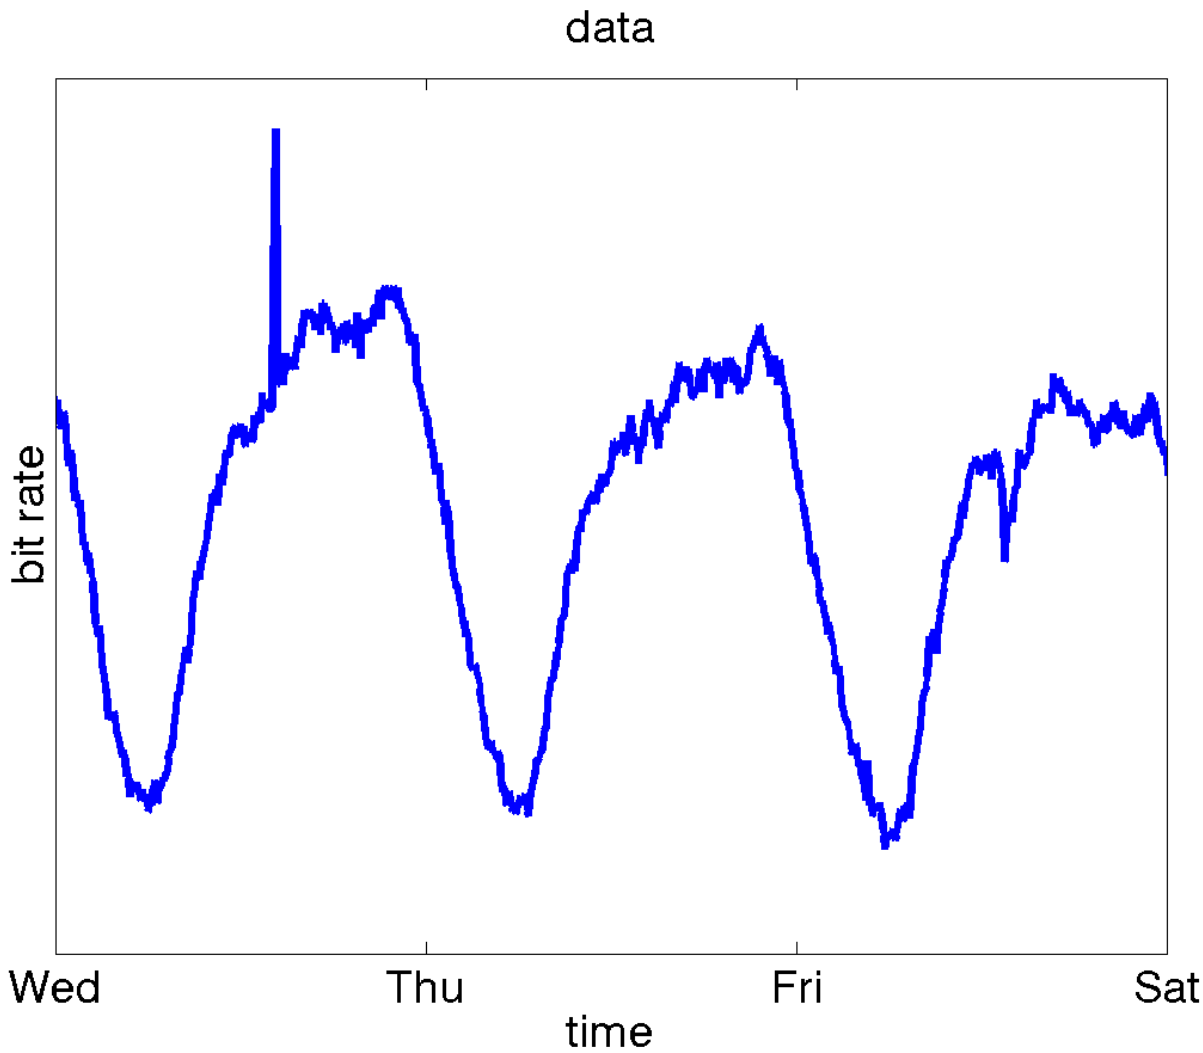

Example: compression (by averaging)

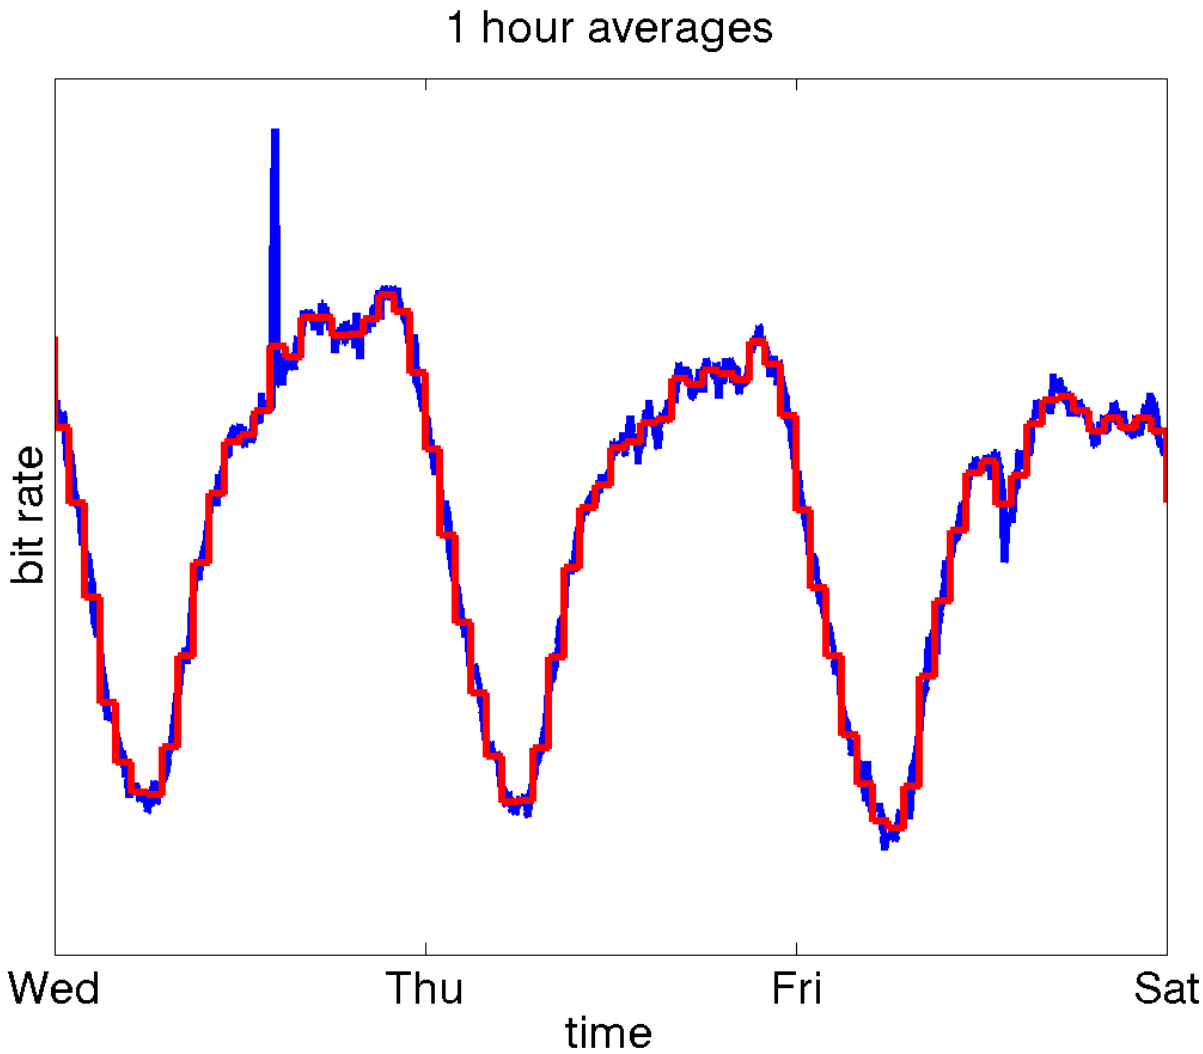

Example: compression (Haar)

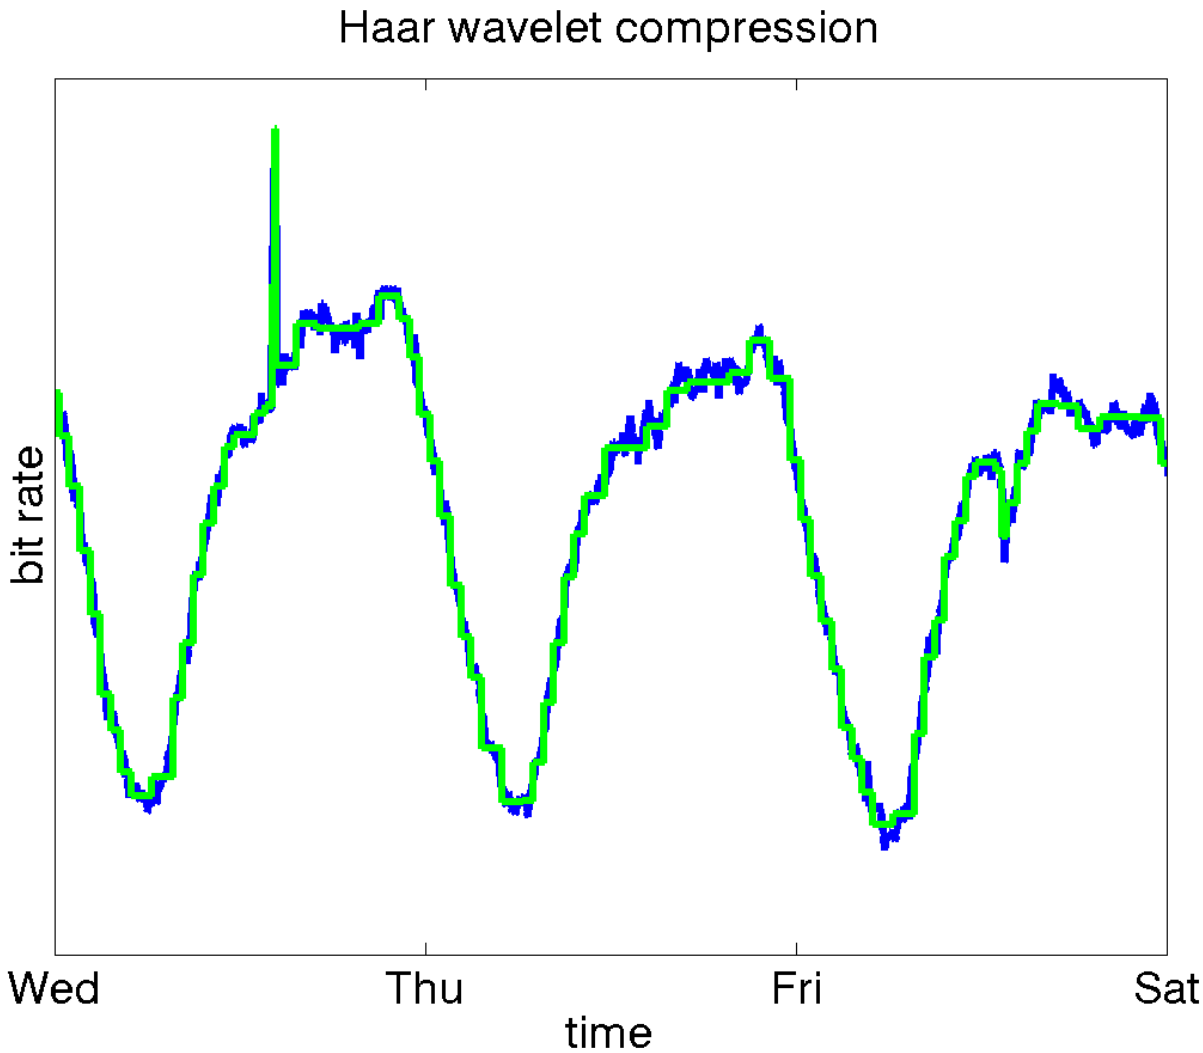

Example: compression (Daubechie's)

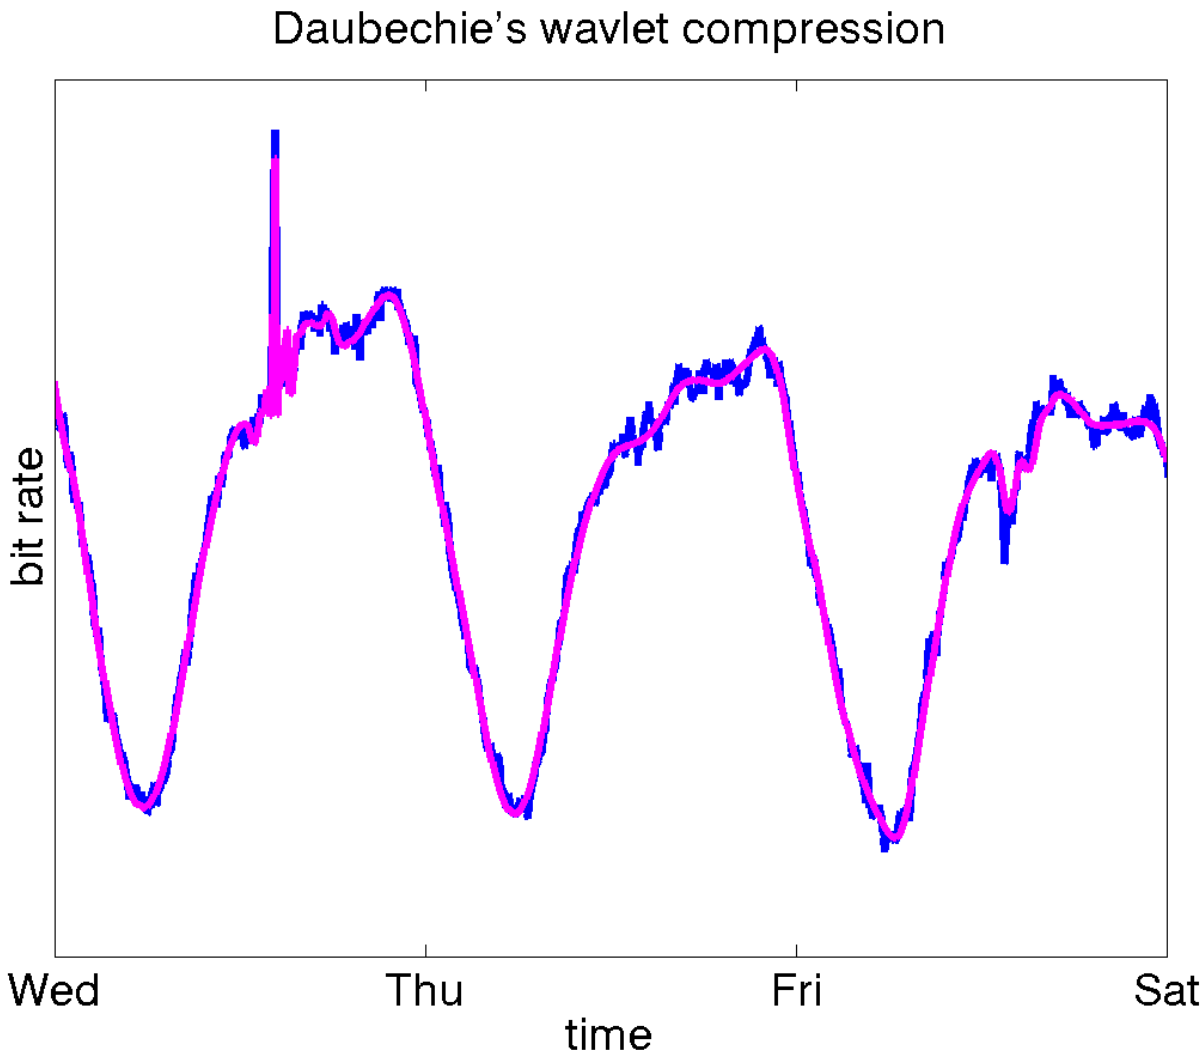

Example: interpolation

- Wavelet based

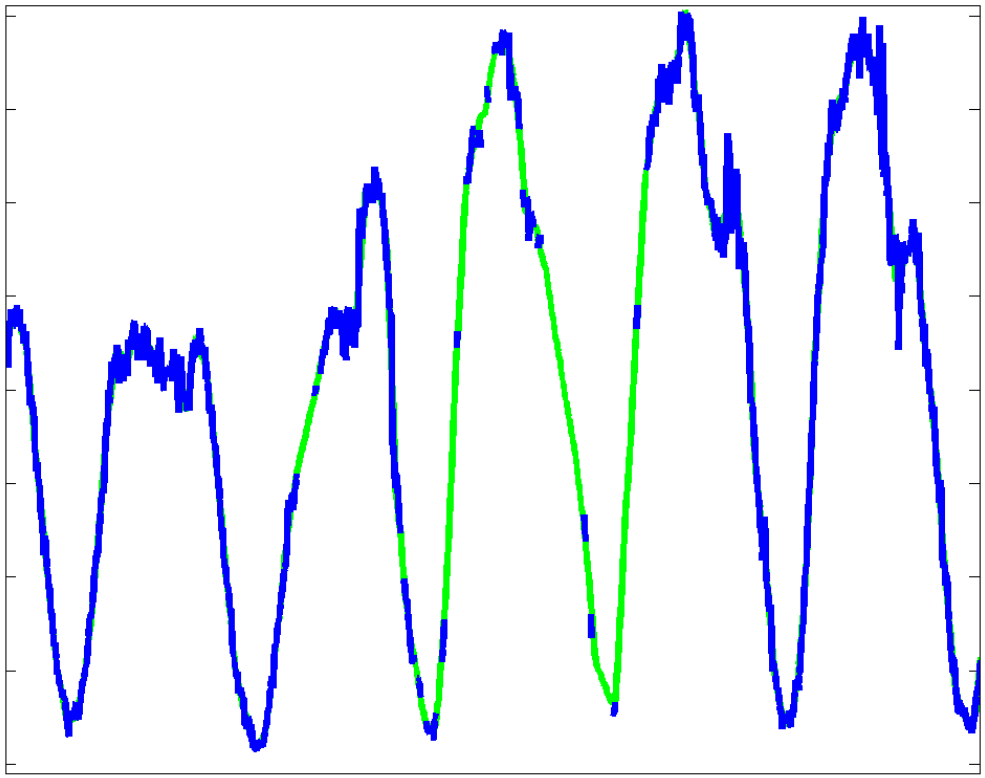

Example: anomaly detection

- Wavelet based

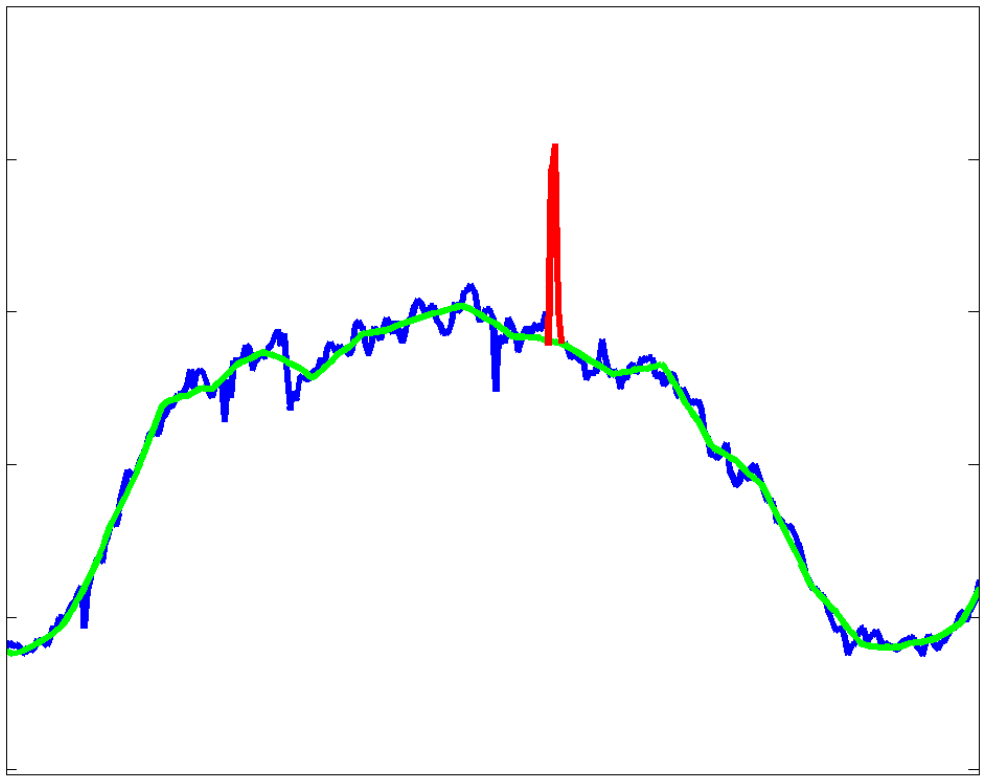

Wavelets, wavelets everywhere and not a …

- Parameter tuning
  - How do know it will work next time?
- Scale of dyadic grid doesn't match patterns in data
  - 5 minute measurements
  - 24 hour cycle, 7 day cycle
  - But dyadic grid is in powers of 2
  - CWT looses many of the advantages of DWT
- Example
  - Compression
  - Look for parameters/wavelet that don't loose important data
  - What is the important data?
- If we had a model it could tell us what is important
  - Compress => estimate model parameters => test difference

Part III: Modeling

- Putting together theory from
  - Time series analysis
  - Traffic theory
- To SNMP data
  - In particular for backbone traffic

Total traffic into a city for 2 weeks

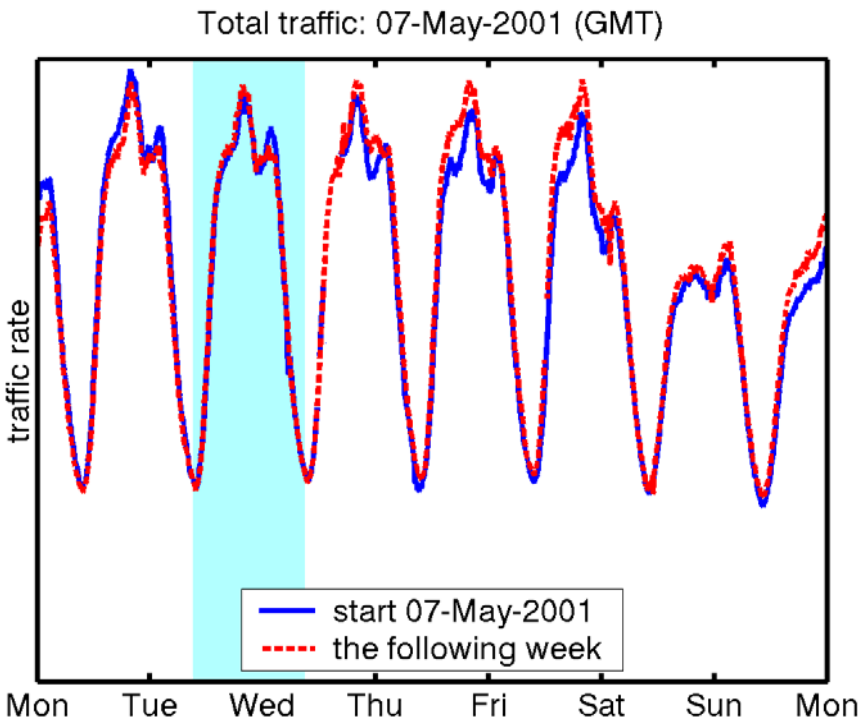

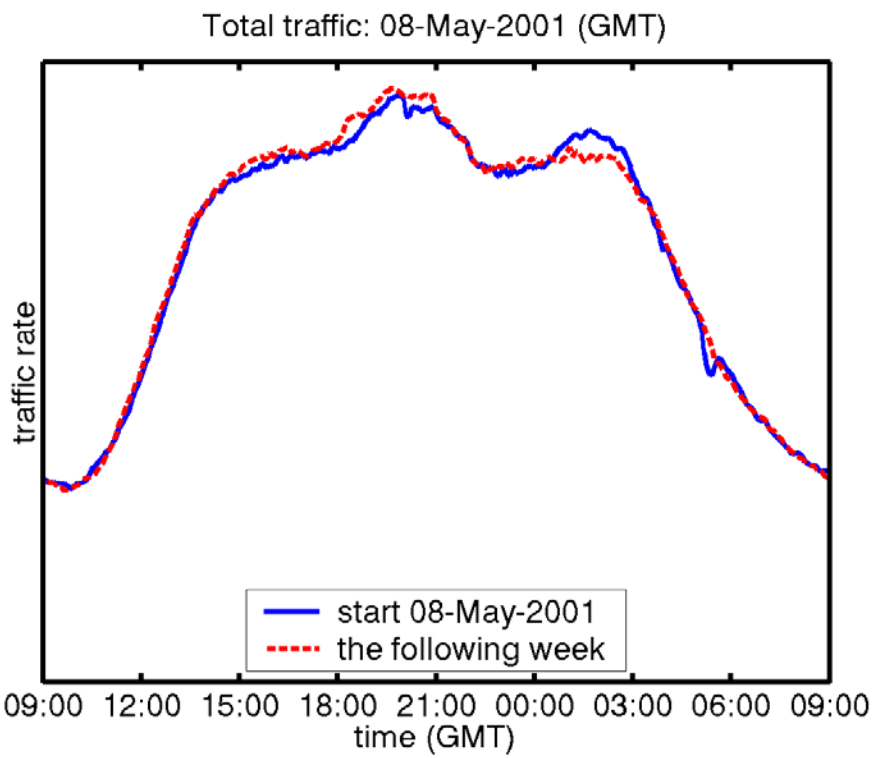

Model

- Traffic data has several components
  - Trend, T_t
    - Long term changes in traffic
  - Seasonal (periodic) component, S_t
    - Daily and weekly cycles
  - Stationary stochastic component, W_t
    - Normal variation
  - Transient anomalies, I_t
    - DoS, Flash crowds, Rerouting (BGP, link failures)
- many ways you could combine these components
  - standard time series analysis
    - Sum $X_t = T_t + S_t + W_t + I_t$
    - Product $X_t = T_t S_t W_t I_t$
    - Box-Cox transform

A Simple Model (for backbone traffic)

- Based on Norros model
- Non-stationary mean
- Stochastic component unspecified (for the moment)

$$x_t = m_t + \sqrt{a m_t}\, W_t + I_t$$

$$m_t = T_t S_t$$

Why this model?

- Behaves as expected under multiplexing

$$x = \sum_i x_i$$

$$m = \sum_i m_i$$

$$a = \frac{\sum_i m_i a_i}{\sum_i m_i}$$

- Good model for backbone traffic
  - Lots of multiplexing
- Simple, estimable parameters, flexible, can make predictions, data supports it

What does a model get you?

- Decomposition
  - MA for trend (window > period of seasonal component)
  - SMA for seasonal component (average at same time of day/week)
  - Several methods for segmenting I_t
- Interpolation
  - Linear, or wavelet based for short gaps (<3 hours)
  - Model based for long gaps (>3 hours)
- Understanding of the effect of multiplexing
  - Should be understood
    - People still seem to misunderstand
  - How smooth is backbone traffic (is it LRD)
- Capacity planning

Example: decomposition

Data => Decomposition

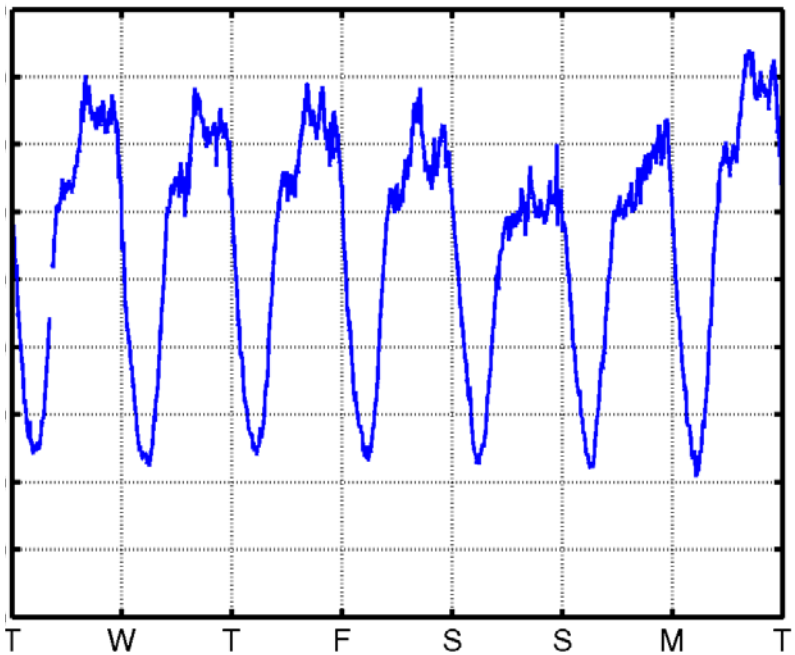

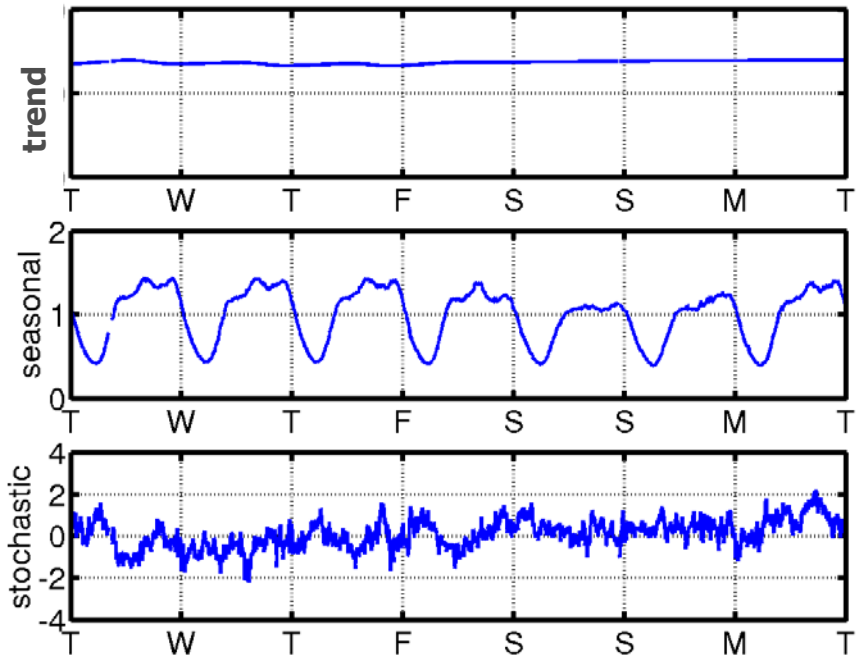

Example: interpolation

- Model based vs linear

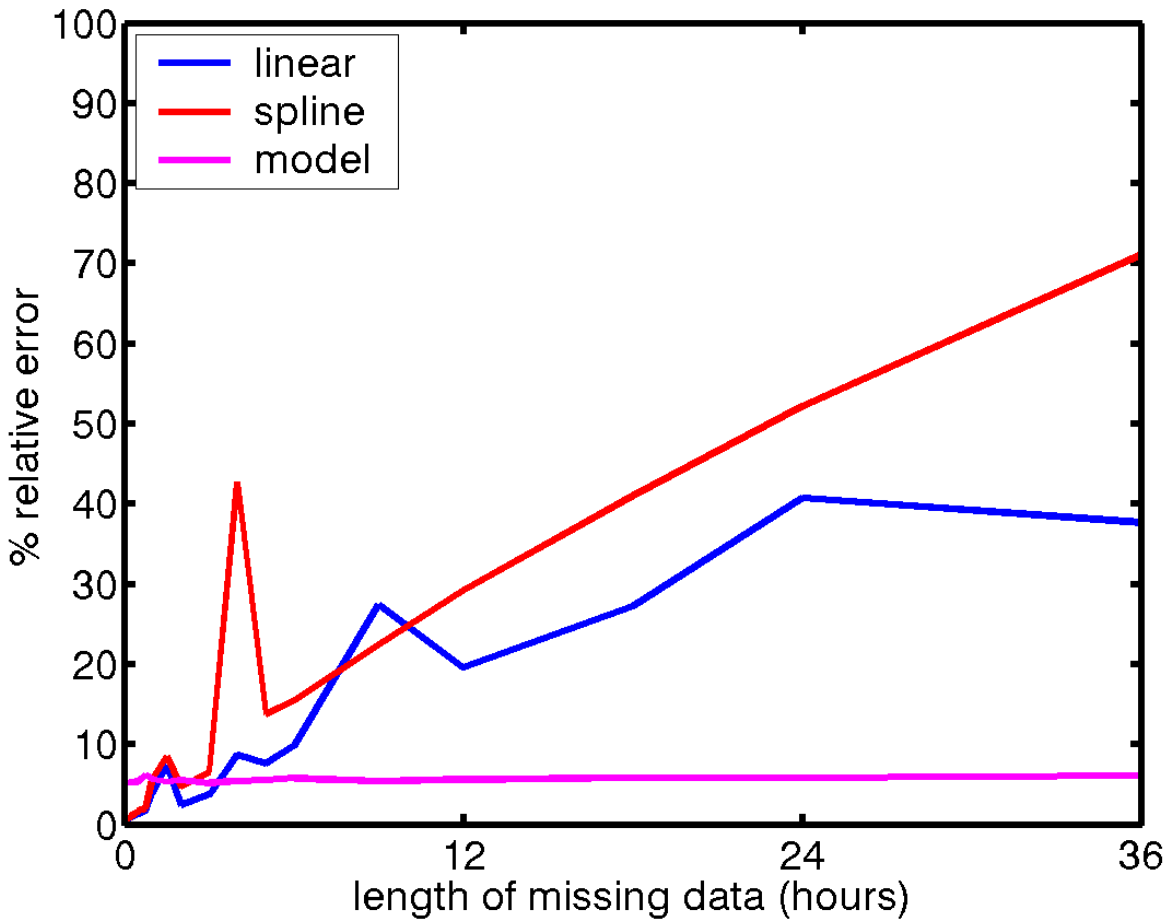

Conclusion

- SNMP is a good data source
  - Available everywhere
  - You need to do some work to extract useful data
  - There is still more info. to get (packet traces, flow data, …)
- Wavelets are a flexible tool for extracting info
  - Not always obvious how to set parameters
- Traffic model gives you a little more
  - A framework for other algorithms
  - A way to decide what information is important
  - A way of seeing how smooth traffic really is
    - Effect of multiplexing
- Algorithms are applicable to other traffic data