ML for Transport Networks: Opportunities and Challenges

George Rouskas

March 5, 2019


cisco
About Cisco / Corporate Development / Acquisitions /
Acquisitions
Singularity NETWORKS
Cisco Completes Acquisition of Singularity Networks, a Network Analytics Company
Singularity Networks analytics capabilities strengthen Cisco Crosswork Network Automation portfolio.
Read announcement
Overview
FAQs
Customers & Partners
Services & Support
SAN JOSE, Calif. - On February 11, 2019, Cisco announced it has completed the acquisition of Singularity Networks, Inc., a privately-held network infrastructure analytics company. The company is based in Denver.
Singularity Networks offers a network monitoring and analytics solution for complex enterprise and service provider networks. The Singularity Networks solution captures, enriches, and analyzes network data to help service providers, web companies, and enterprises improve network performance, manage costs, gain visibility, and reduce downtime.
Cisco will integrate Singularity Networks analytics platform capabilities into its Crosswork Network Automation

Branches of Machine Learning

- **Supervised Learning**
  - Training data used to learn a function to map input to output
- **Unsupervised learning**
  - Learn pattern from input without training or explicit output
- **Reinforcement Learning**
  - No training, learning through feedback (positive/negative) from environment

Classification

- Prediction of discrete output values

Regression

- Prediction of continuous output values

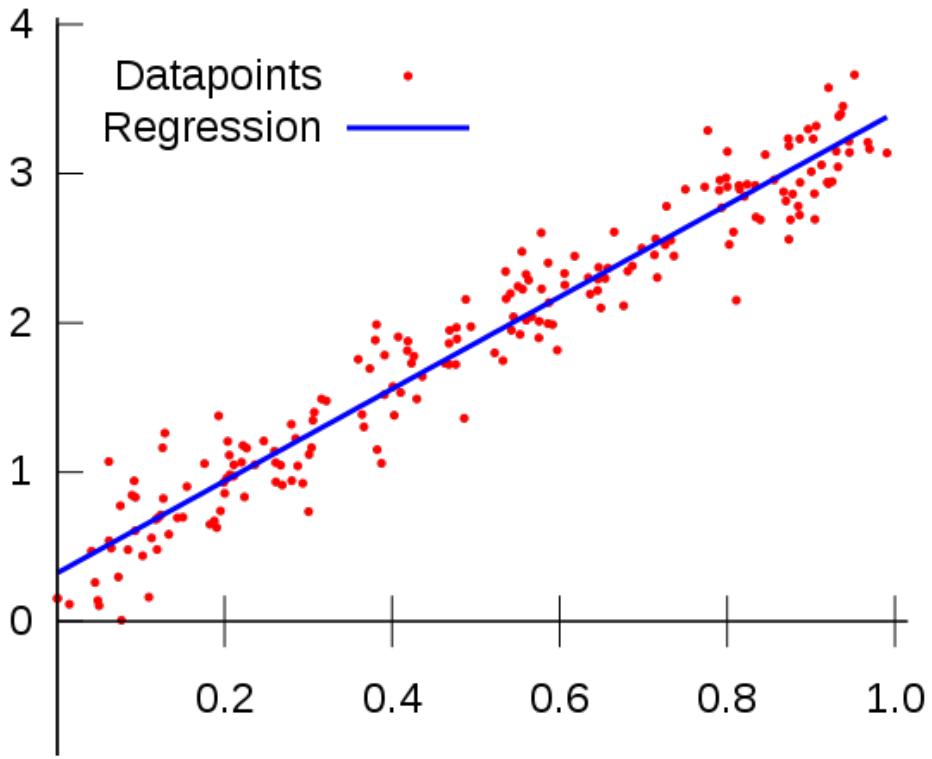

Clustering

- Partition of dataset into groups

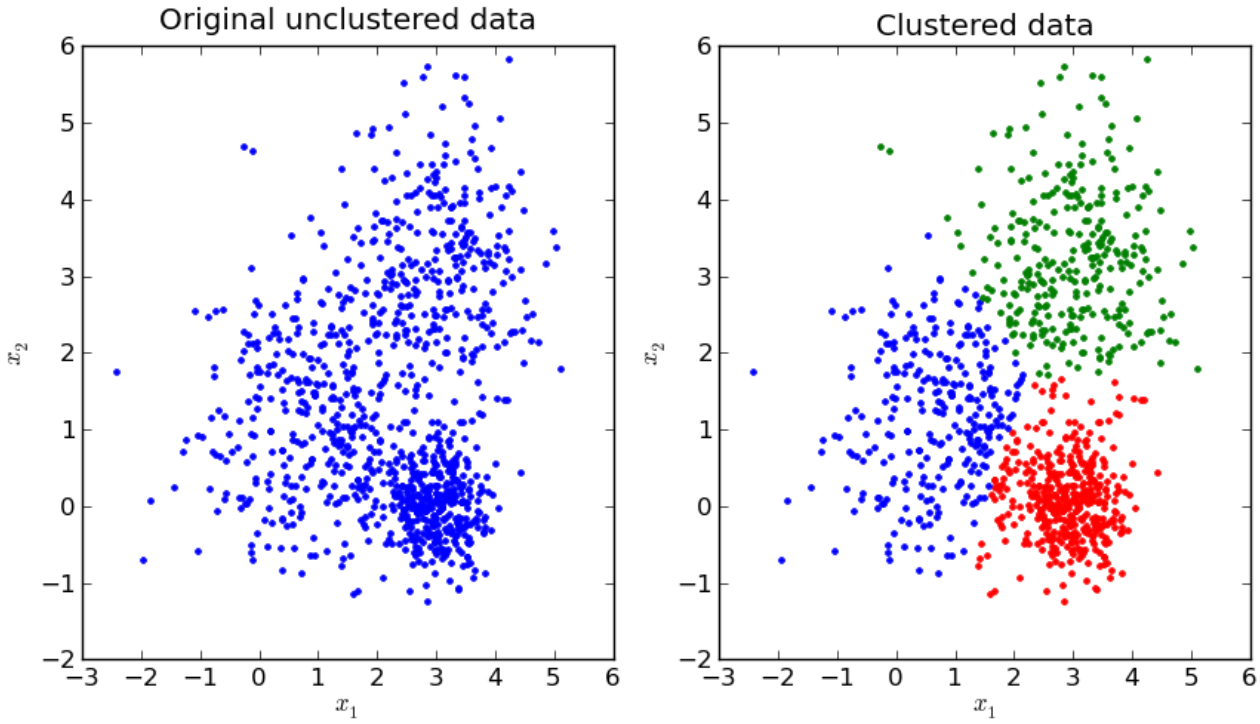

Control

- Control of nonlinear systems

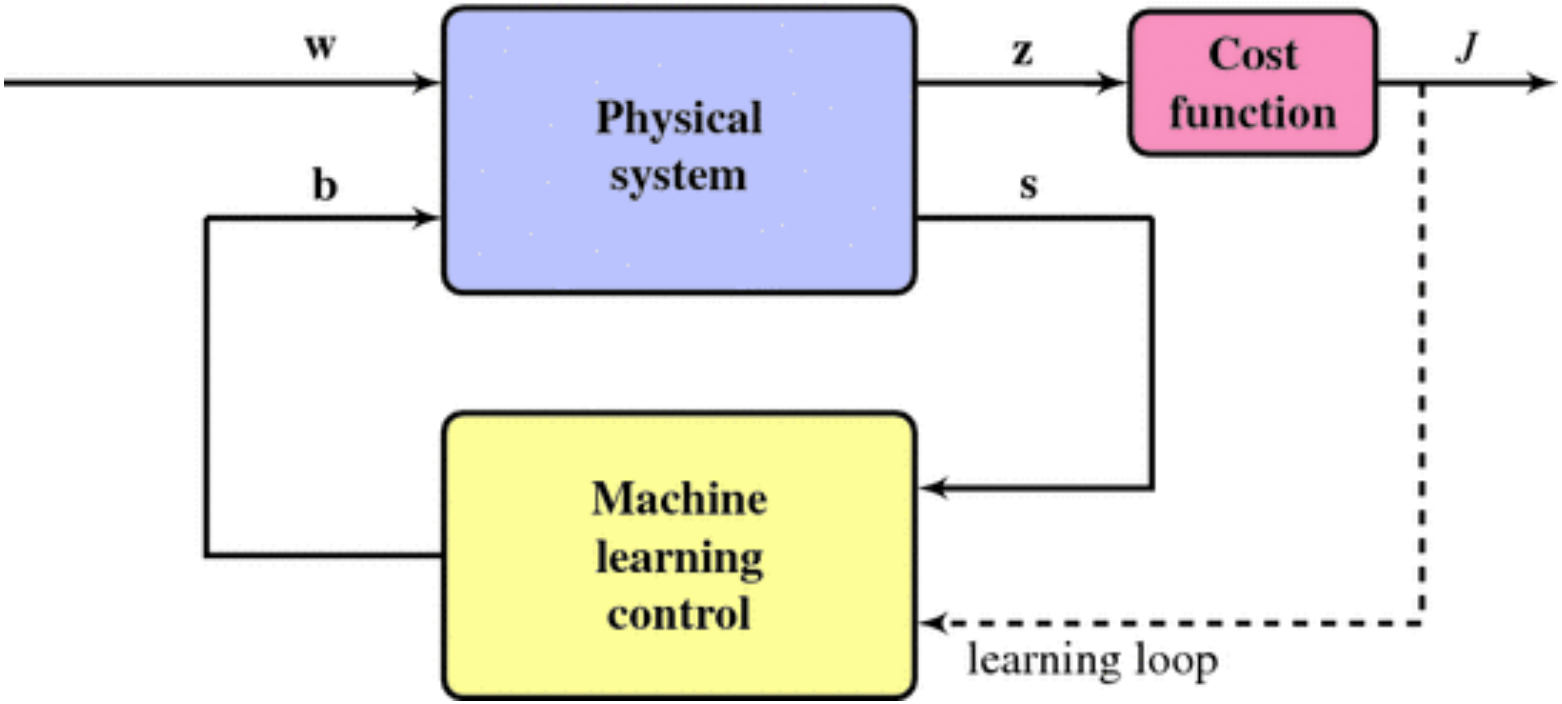

1650 JOURNAL OF LIGHTWAVE TECHNOLOGY, VOL. 36, NO. 9, MAY 1, 2018

Lightpath Management in SDN-Based Elastic Optical Networks With Power Consumption Considerations

Yu Xiong, Jin Shi, Yaya Yang, Yi Lv, and George N. Rouskas

***Abstract*—Elastic optical networks (EONs) are increasingly used to interconnect data centers. As interdatacenter traffic is expected to continue to grow at high rates for the foreseeable future, it becomes imperative to address the respective growth in power consumption within EONs. In this paper, we take the first steps toward enhancing the energy efficiency of interdatacenter EONs. Specifically, we use traffic prediction techniques for centralized lightpath management that leverages the capabilities of software defined networking (SDN) technology. Our objective is to eliminate unnecessary lightpath termination and re-establishment operations, and in turn decrease switching power consumption within the network. Our evaluation study confirms that the proposed algorithm is effective in achieving substantial savings in power consumption while maintaining a bandwidth blocking ratio at levels comparable to those of earlier algorithms.**

***Index Terms*—Elastic optical networks, energy efficiency, lightpath management, software defined networking, switching power, traffic prediction.**

necessitate an increase in the scale of DC facilities, which in turn leads to higher requirements in terms of inter-DC communication to support data backup, data synchronization, and collaboration between different DCs. The demands placed on inter-DC communication call for appropriate network technologies to interconnect DCs effectively and efficiently.

Elastic optical networks (EONs) [1], [2] are widely regarded as the most promising technology for interconnecting DCs, and have been studied extensively. EONs utilize bandwidth variable optical transponders (BV-OT) and bandwidth variable optical cross-connects (BV-OXC) that operate on a set of spectrally-contiguous frequency slots to set up lightpaths. Since these frequency slots occupy a much narrower bandwidth than the conventional wavelength channels, EONs can provision bandwidth adaptively according to actual traffic demands [3], and hence may meet the requirements of DC traffic. At the same time, the technological heterogeneity and resource diversity be-

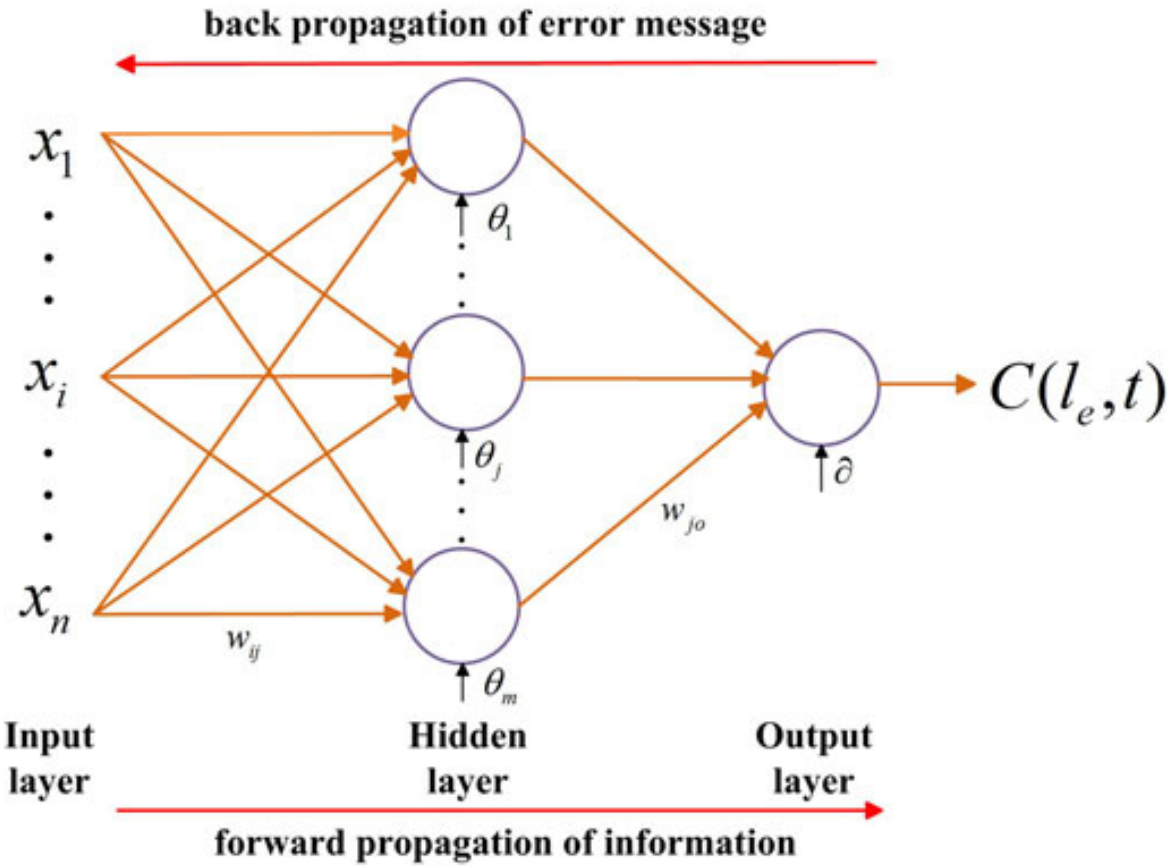
back propagation of error message
x_1
x_i
x_n
θ_1
θ_j
θ_m
w_{ij}
w_{jo}
∂
$C(l_e,t)$
Input layer
Hidden layer
Output layer
forward propagation of information

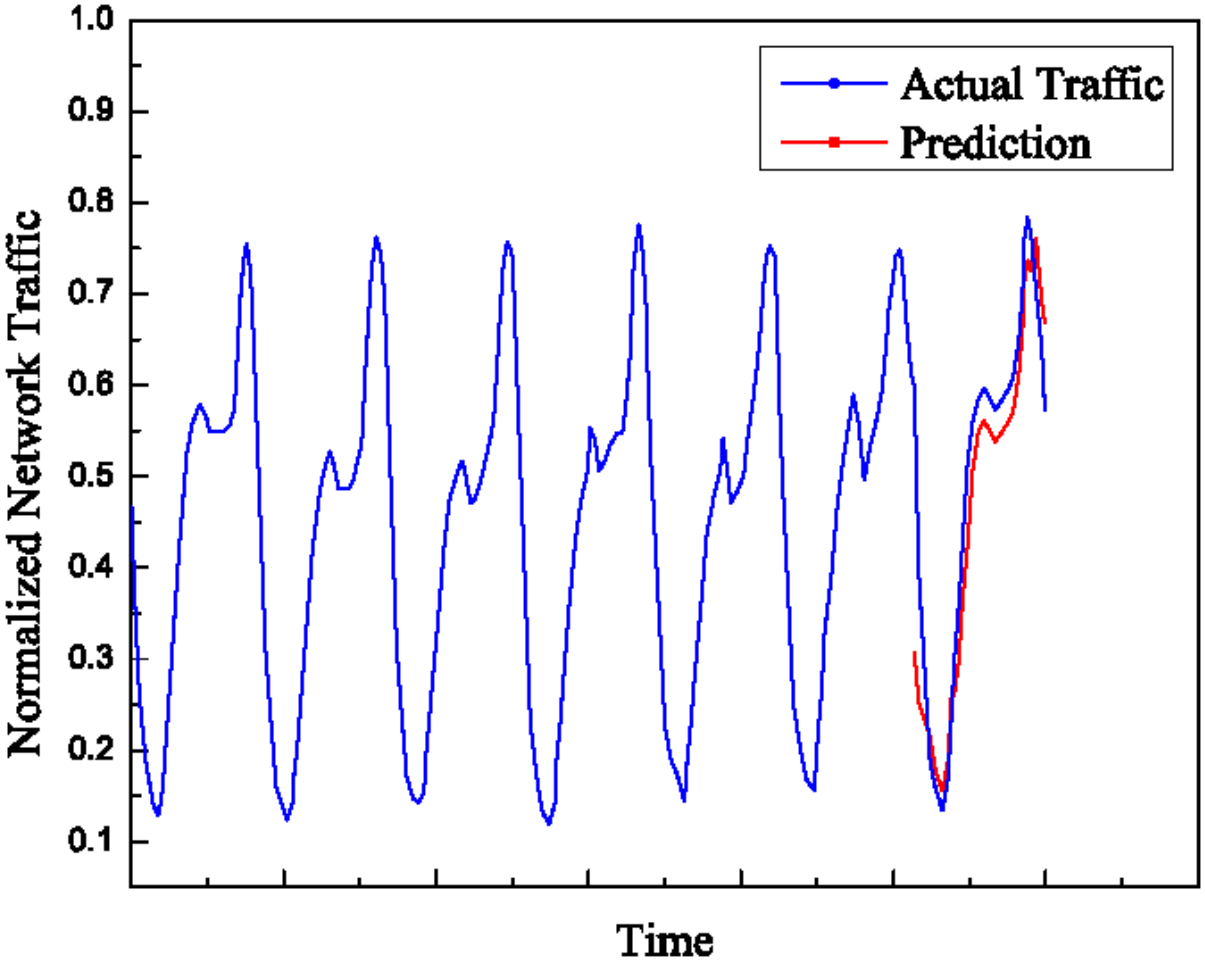
Actual Traffic
Prediction
Normalized Network Traffic
Time
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1

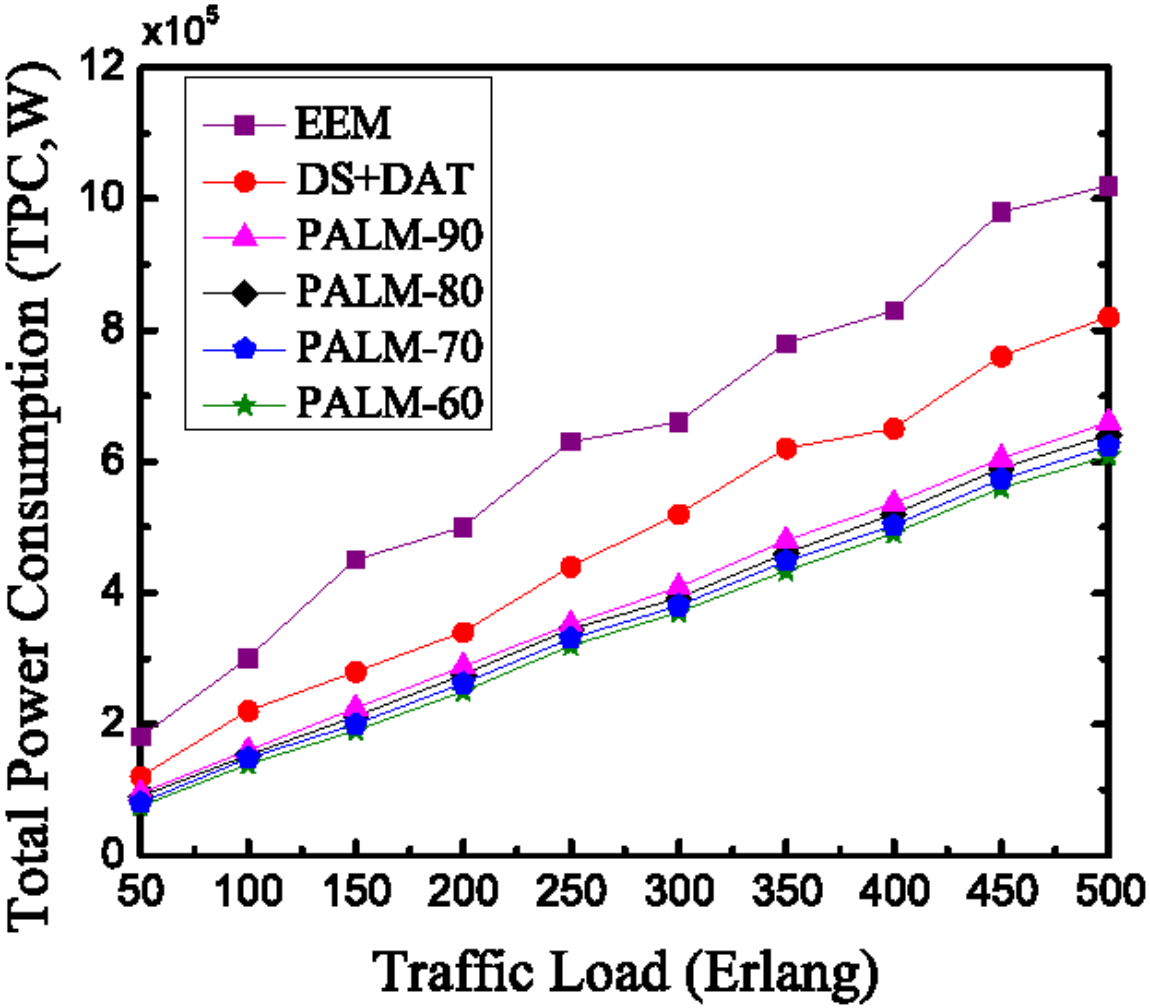
x10^5
12
10
8
6
4
2
0
Total Power Consumption (TPC, W)
EEM
DS+DAT
PALM-90
PALM-80
PALM-70
PALM-60
50
100
150
200
250
300
350
400
450
500
Traffic Load (Erlang)

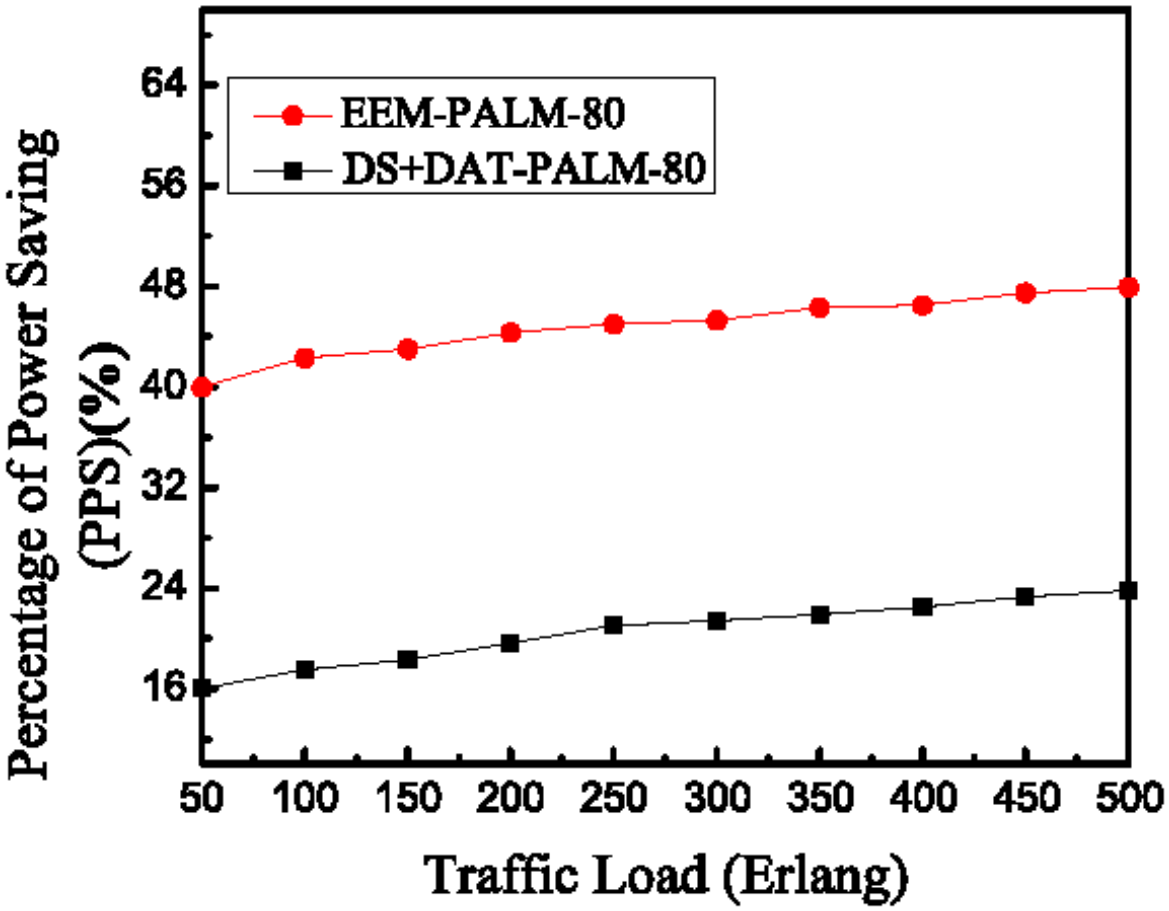
EEM-PALM-80
DS+DAT-PALM-80
Percentage of Power Saving (PPS)(%)
64
56
48
40
32
24
16
50
100
150
200
250
300
350
400
450
500
Traffic Load (Erlang)

Opportunities

- Transport networks: heterogeneous, time-varying, multi-layer, diverse, packet/optical
  - Difficult to model analytically
- Lots of data: logs, monitors, signaling, configuration,…
- SDN: large decision space and set of tuning params
- Interaction with 5G networks, many monitoring points
  - Highly complex/dynamic optimization problems

Applications

- Quality of Transmission (QoT) estimation
- Determination of optimal modulation format
- Predictive repair/performance strategies
- Proactive reconfiguration and autonomic network mgmt
- Attack and intrusion detection
- Concerted operation of
  - multiple network layers
  - networks/compute resources
- …

Challenges

- Bias in modeling process and feedback loops
- Black-box approach: “what” but not “why”
- Network hardware not yet open/interoperable
- Slow deployment of SDN in transport networks
- Low quality of available data