SWINBURNE UNIVERSITY OF TECHNOLOGY

Statistical Traffic Classification

Jason But

jbut@swin.edu.au

Centre for Advanced Internet Architectures (CAIA)
Swinburne University of Technology

Performance

Statiscal Classification

- Originally demonstrated with real-time classification of game traffic
- Expanded to Skype
- Then BitTorrent
- Performance is impressive – $\geq$ 95% Recall and Precision

How does it Work?

- Minimal state
  - Maintain counters per sub-flow
  - Calculate features
- Stateless classification

Rapid Flow Classification

Classification processSubflow Classification

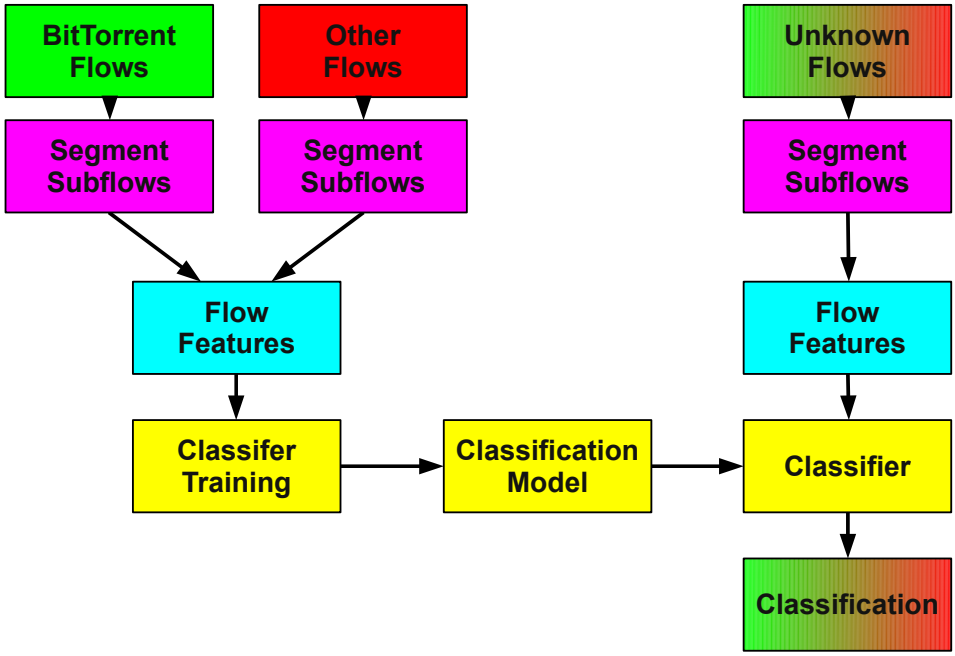

- We train the classifier to detect sub-flows

Improving Performance

- Can we improve performance by using more state?
- Using a feedback loop of old(er) classifications

Applying an Output Filter

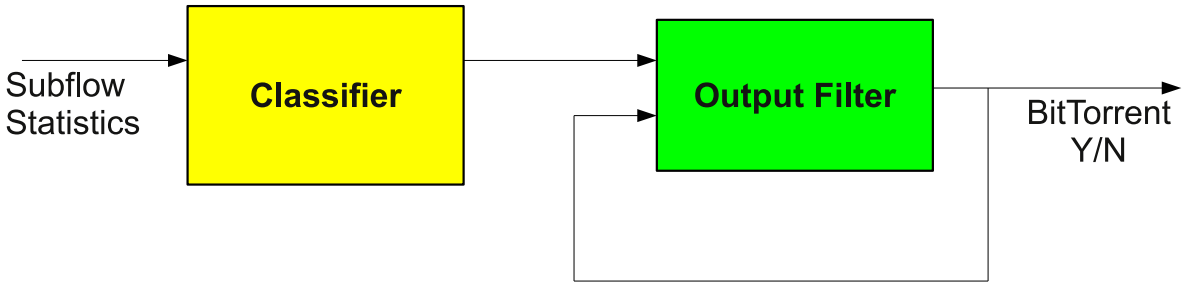

Improving Performance

Tradeoffs

- Increased memory consumption
  - How many previous classifications?
  - How to calculate filtered output
- Increased CPU resources

Is it worth it?

- Smaller sub-flow sizes
- Faster classification

Other Considerations

- Testing with encrypted BitTorrent
- More generic peer-2-peer classification