

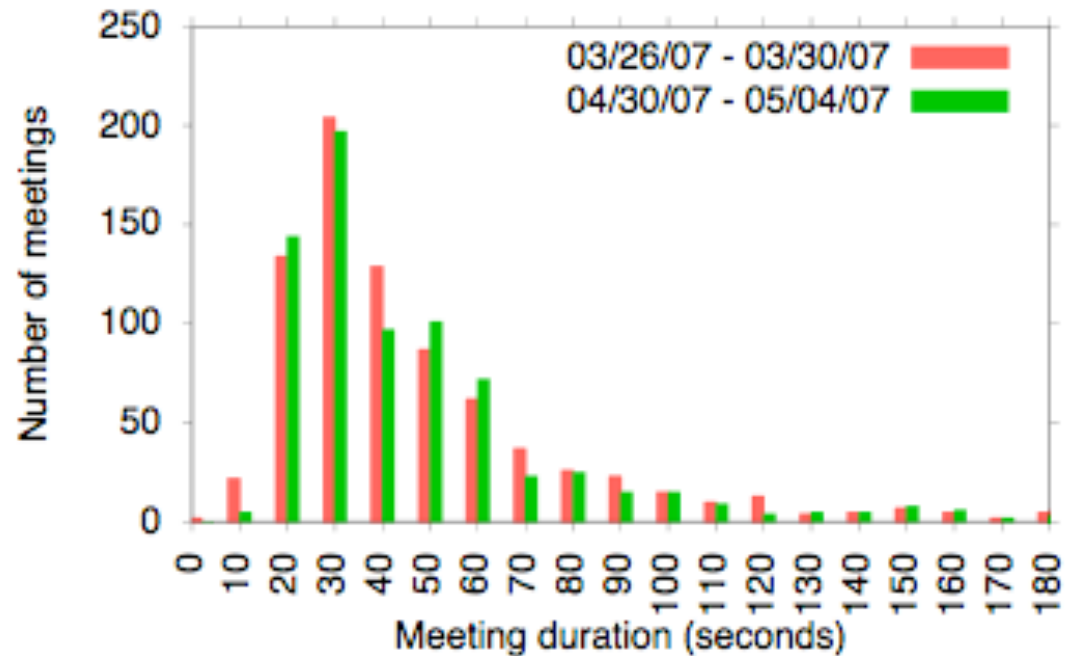

# Web Search From a Bus

**Brian Neil Levine**

Supported in part by  
**NSF-0133055**  
**CNS-0519881**




# DieselNet



- 40 buses equipped with wireless networking
- APs met about every 6 minutes
- <http://prisms.cs.umass.edu/dome>

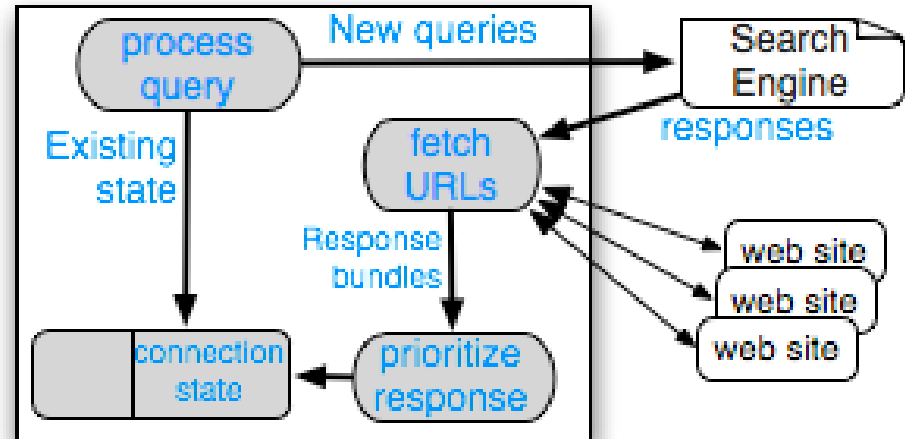
# Web Search

- Web search and retrieval is typically an interactive process in a tethered network.
  - don't work with intermittent connectivity


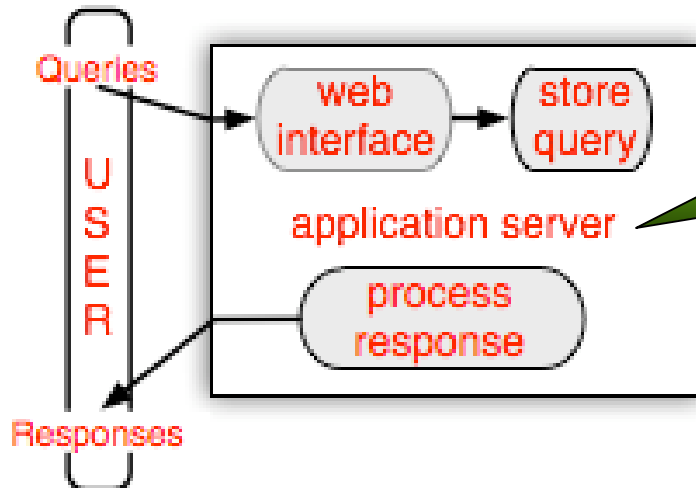


# Thedu Architecture

## Server-side Proxy



## Mobile Node

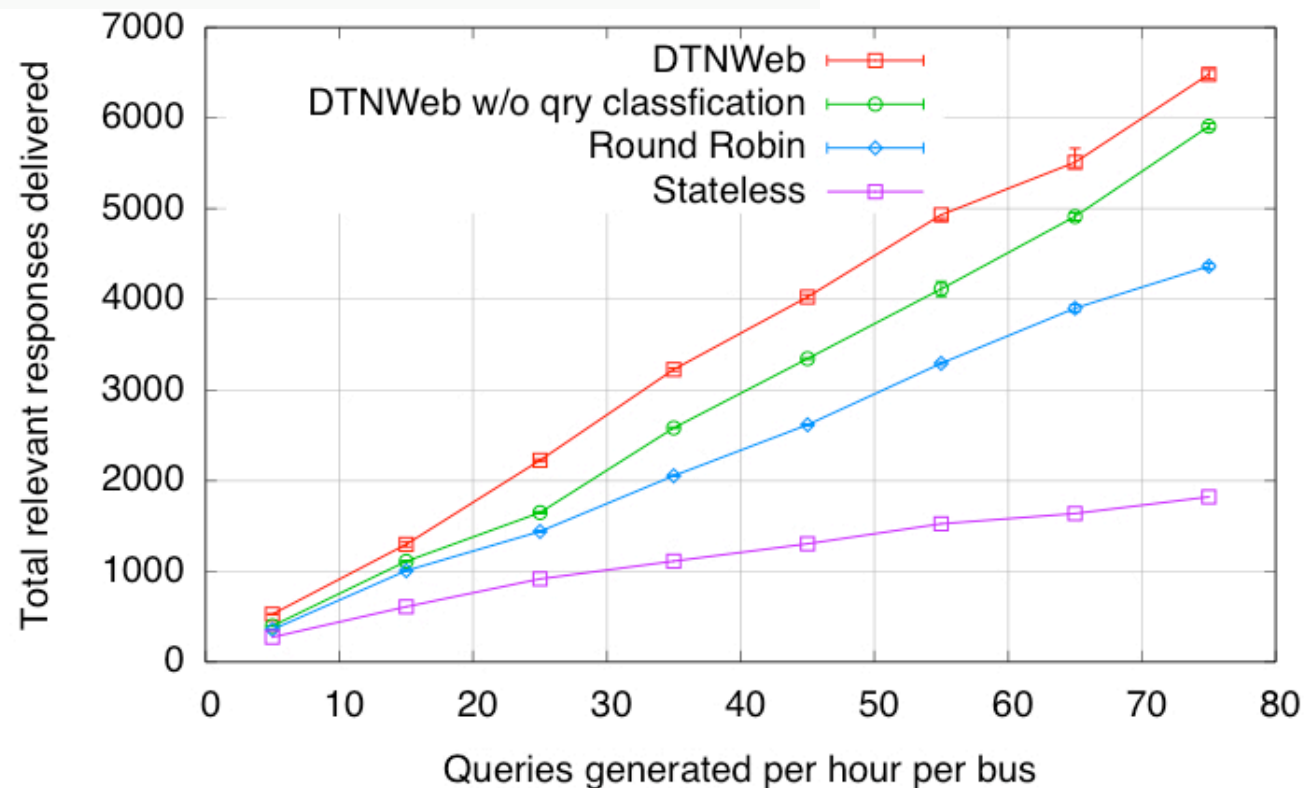


# Result Prioritization

- Relevance scores are normalized across disparate queries
  - Scores then converted into a probability
  - Queries are classified as *homepage* queries (or not) — requires one result.
- Sort by relevance score
  - IF homepage query, send best result incrementing  $E(q) += P(r)$ ; stop if  $E(q) > 1$
  - ELSE, send result.

# Deployment and Evaluation

Statistic	Thedu Proxy	No Proxy
Queries:	780	743
Web pages:	5639	1207
Avg resps. per query:	7.2	1.6
Relevant web pages:	1630	401
Queries with at least 1 relevant web page:	529 (68%)	291 (39%)



# Conclusion

- Thedu: web search and retrieval in an intermittently connected environment
- Interactive web search converted into prefetching transaction
- Deployed for one week in DieselNet
  - First relevant response received within 3 min
  - Performance will improve in urban environments
- *This work will appear in ACM CHANTS 2007*  
Co-authors: Aruna Balasubramanian, Yun Zhou, W. Bruce Croft, and Arun Venkataramani
- See also Sigcomm 2007, Mobicom 2007, Mobihoc 2007.