On the validity (or otherwise) of IEEE 802.11 mathematical modeling hypotheses

Ken Duffy

Joint work with Kaidi Huang and David Malone

Hamilton Institute, National University of Ireland Maynooth

Swinburne University of Technology, May 14^{th} 2009

Talk outline.

Talk outline.

- The IEEE 802.11 CSMA/CA MAC.

Talk outline.

- The IEEE 802.11 CSMA/CA MAC.
- Recent advances in mathematical modeling.

Talk outline.

- The IEEE 802.11 CSMA/CA MAC.
- Recent advances in mathematical modeling.
- Implicit approximations made to enable analytic tractability.

Talk outline.

- The IEEE 802.11 CSMA/CA MAC.
- Recent advances in mathematical modeling.
- Implicit approximations made to enable analytic tractability.

- Directly testing these hypotheses with test-bed data.

Talk outline.

- The IEEE 802.11 CSMA/CA MAC.
- Recent advances in mathematical modeling.
- Implicit approximations made to enable analytic tractability.
- Directly testing these hypotheses with test-bed data.
- Summary, an epilogue and conclusions.

The 802.11 DCF

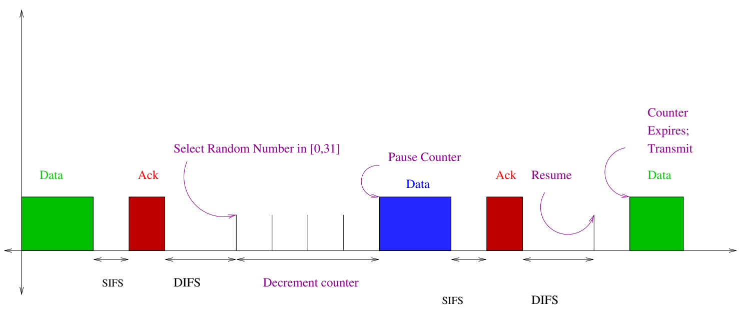

Figure: 802.11 MAC operation

The 802.11 MAC flow diagram

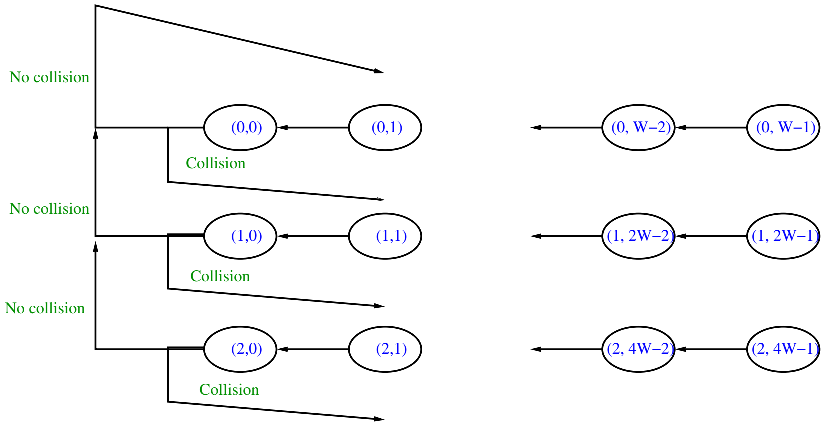

Figure: Saturated 802.11 MAC operation

Popular mathematical modeling approaches

- P-persistent:approximate the back-off distribution be a geometric with the same mean. E.g. work by Marco Conti and co-authors (F Cali, M Conti, E Gregori, P Aleph IEEE/ACM ToN 2000).

Popular mathematical modeling approaches

- P-persistent:approximate the back-off distribution be a geometric with the same mean. E.g. work by Marco Conti and co-authors (F Cali, M Conti, E Gregori, P Aleph IEEE/ACM ToN 2000).
- Mean-field Markov models: seminal work by Bianchi (IEEE Comms L. 1998, IEEE JSAC 2000).

Bianchi's approach

Observation: each individual station's impact on overall network access is small.

Bianchi's approach

Observation: each individual station's impact on overall network access is small.
Mean field approximation: assume a fixed probability of collision at each attempted transmission p, irrespective of the past.

Bianchi's approach

Observation: each individual station's impact on overall network access is small.
Mean field approximation: assume a fixed probability of collision at each attempted transmission p, irrespective of the past.

Each station's back-off counter then a Markov chain.

Mean-field Markov Model's Chain

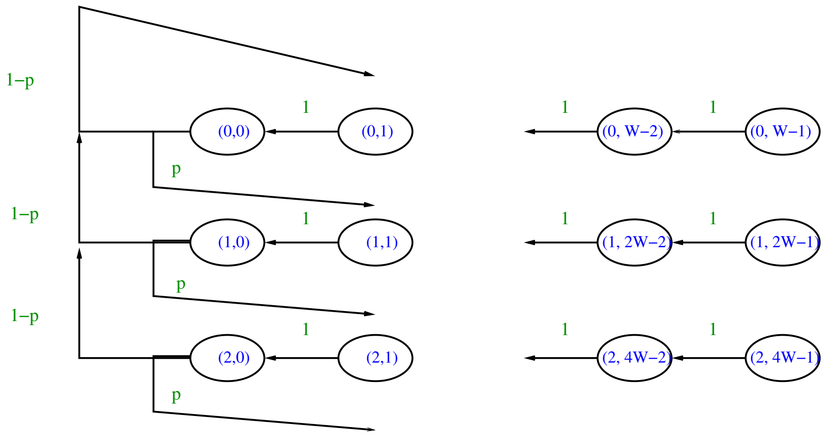

Figure: Individual's Markov Chain if p known

Mean-field Markov Overview

Stationary distribution gives the probability the station attempts transmission in a typical slot

Mean-field Markov Overview

Stationary distribution gives the probability the station attempts transmission in a typical slot

$$\tau(p) = \frac{2(1-2p)}{(1-2p)(W+1) + pW(1-(2p)^m)}.$$

Mean-field Markov Overview

Stationary distribution gives the probability the station attempts transmission in a typical slot

$$\tau(p) = \frac{2(1-2p)}{(1-2p)(W+1) + pW(1-(2p)^m)}.$$

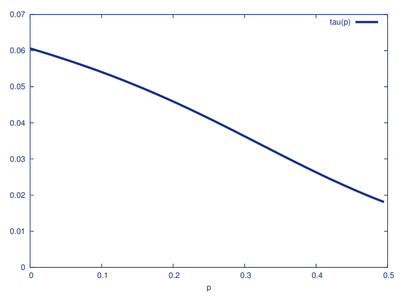

Figure: Attempt probability $\tau(p)$ vs p

The self-consistent equation

Network of N stations.

The self-consistent equation

Network of N stations. Mean field decoupling idea:

The self-consistent equation

Network of N stations. Mean field decoupling idea: the impact of **every** station on the network access of the others is small,

The self-consistent equation

Network of N stations. Mean field decoupling idea: the impact of **every** station on the network access of the others is small, so that

$$1 - p = (1 - \tau(p))^{N-1}. \quad (1)$$

Solution of equation (1) determines the network's "real" p^*.

The self-consistent equation

Network of N stations. Mean field decoupling idea: the impact of **every** station on the network access of the others is small, so that

$$1 - p = (1 - \tau(p))^{N-1}. \quad (1)$$

Solution of equation (1) determines the network's "real" p^*.

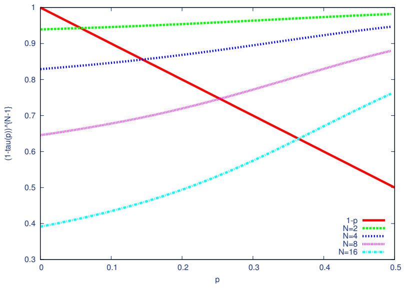

Figure: $1 - p$ and $(1 - \tau(p))^N$ for $N = 2, 4, 8 \,\&16$

Example developments

- Unsaturated 802.11, Small buffer: Ahn, Campbell, Veres and Sun, IEEE Trans. Mob. Comp., 2002; Ergen, Varaiya, ACM-Kluwer MONET, 2005; Malone, K.D., Leith, IEEE/ACM Trans. Network., 2007.

Example developments

- Unsaturated 802.11, Small buffer: Ahn, Campbell, Veres and Sun, IEEE Trans. Mob. Comp., 2002; Ergen, Varaiya, ACM-Kluwer MONET, 2005; Malone, K.D., Leith, IEEE/ACM Trans. Network., 2007.
- Unsaturated 802.11, Big buffer: Cantieni, Ni, Barakat and Turletti, Comp. Comm., 2005; Park, Han and Ahn, Telecomm. Sys., 2006; K.D. and Ganesh, IEEE Comm. Lett., 2007.

Example developments

- Unsaturated 802.11, Small buffer: Ahn, Campbell, Veres and Sun, IEEE Trans. Mob. Comp., 2002; Ergen, Varaiya, ACM-Kluwer MONET, 2005; Malone, K.D., Leith, IEEE/ACM Trans. Network., 2007.
- Unsaturated 802.11, Big buffer: Cantieni, Ni, Barakat and Turletti, Comp. Comm., 2005; Park, Han and Ahn, Telecomm. Sys., 2006; K.D. and Ganesh, IEEE Comm. Lett., 2007.
- 802.11e, Saturated: Kong, Tsang, Bensaou and Gao, IEEE JSAC, 2004; Robinson and Randhawa, IEEE JSAC, 2004. Unsaturated: Zhai, Kwon and Fang, WCMC, 2004. Chen, Xhai, Tian and Fang, IEEE Trans. W. Commun., 2006.

Example developments

- Unsaturated 802.11, Small buffer: Ahn, Campbell, Veres and Sun, IEEE Trans. Mob. Comp., 2002; Ergen, Varaiya, ACM-Kluwer MONET, 2005; Malone, K.D., Leith, IEEE/ACM Trans. Network., 2007.
- Unsaturated 802.11, Big buffer: Cantieni, Ni, Barakat and Turletti, Comp. Comm., 2005; Park, Han and Ahn, Telecomm. Sys., 2006; K.D. and Ganesh, IEEE Comm. Lett., 2007.
- 802.11e, Saturated: Kong, Tsang, Bensaou and Gao, IEEE JSAC, 2004; Robinson and Randhawa, IEEE JSAC, 2004. Unsaturated: Zhai, Kwon and Fang, WCMC, 2004. Chen, Xhai, Tian and Fang, IEEE Trans. W. Commun., 2006.
- 802.11s, unsaturated: K.D., Leith, Li and Malone, IEEE Comm. Lett., 2006.

Standard approach to model verification

ASK: Do the model throughput and delay predictions match well with results from simulated system?

Standard approach to model verification

ASK: Do the model throughput and delay predictions match well with results from simulated system?
NOT: Make the approximations explicit hypotheses and check them directly.

A warning from hydrology

"The modelling technology has far outstripped the level of our understanding of the physical processes being modeled. Making use of this technology then requires that the gaps in the factual knowledge be filled with assumptions which, although often appearing logical, have not been verified and may sometimes be wrong".

Vit Klemes, WCP-98, WHO, 1985.

Test bed

Figure: PC as AP, 1 PC and 9 PC-based Soekris Engineering net4801 as clients. All with Atheros AR5215 802.11b/g PCI cards. Modified MADWiFi wireless driver for fixed 11 Mbps transmissions and specified queue-size.

What are the hypotheses?

All models:

- $C_k = 1$ if k^{th} transmission results in collision.

What are the hypotheses?

All models:

- $C_k = 1$ if k^{th} transmission results in collision.
- $C_k = 0$ if k^{th} transmission results in success.

What are the hypotheses?

All models:

- $C_k = 1$ if k^{th} transmission results in collision.
- $C_k = 0$ if k^{th} transmission results in success.

Assumptions:

- (A1) $\{C_k\}$ is an independent sequence;

What are the hypotheses?

All models:

- $C_k = 1$ if k^{th} transmission results in collision.
- $C_k = 0$ if k^{th} transmission results in success.

Assumptions:

- (A1) $\{C_k\}$ is an independent sequence;
- (A2) $\{C_k\}$ are identically distributed with $P(C_k = 1) = p$.

Testing (A1): $\{C_k\}$ independent

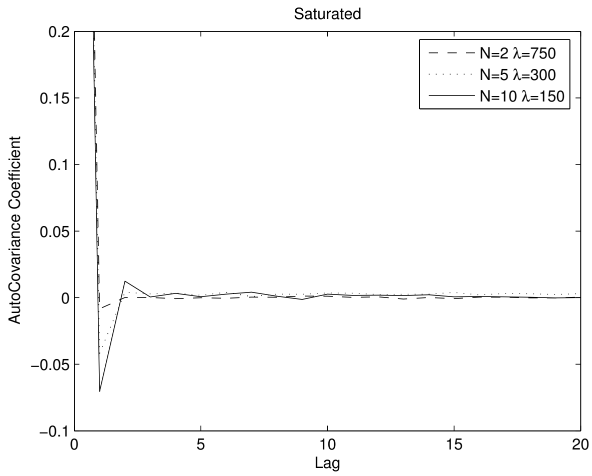

Figure: Saturated $C_1, \ldots, C_K$ normalized auto-covariances. Experimental data, $N = 2, 5, 10$, $K = 2500k, 1200k, 711k$.

Testing (A1): $\{C_k\}$ independent

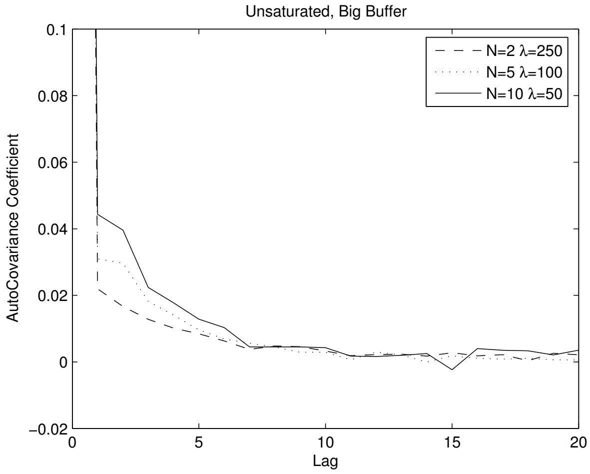

Figure: Unsaturated, big buffer $C_1, \ldots, C_K$ normalized auto-covariances. Experimental data, $N = 2, 5, 10$, $K = 1800k, 750k, 380k$.

Testing (A2): $\{C_k\}$ identically distributed

Record the backoff stage at which the attempt was made.
Probability p_i of collision given backoff stage i.

Testing (A2): $\{C_k\}$ identically distributed

Record the backoff stage at which the attempt was made.
Probability p_i of collision given backoff stage i.
Assumption (A2): $p_i = p$ for all i.

Testing (A2): $\{C_k\}$ identically distributed

Record the backoff stage at which the attempt was made.
Probability p_i of collision given backoff stage i.
Assumption (A2): $p_i = p$ for all i.
MLE

$$\hat{p}_i = \frac{\#\text{collisions at back-off stage } i}{\#\text{transmissions at back-off stage } i}.$$

Testing (A2): $\{C_k\}$ identically distributed

Record the backoff stage at which the attempt was made.
Probability p_i of collision given backoff stage i.
Assumption (A2): $p_i = p$ for all i.
MLE

$$\hat{p}_i = \frac{\#\text{collisions at back-off stage } i}{\#\text{transmissions at back-off stage } i}.$$

Hoeffding's inequality (1963):

$$P(|\hat{p}_i - p_i| > x) \leq 2\exp\left(-2x(\#\text{transmissions at back-off stage } i)\right).$$

To have 95% confidence that $|\hat{p}_i - p_i| \leq 0.01$ requires 185 attempted transmissions at backoff stage i .

Testing (A2): $\{C_k\}$ identically distributed

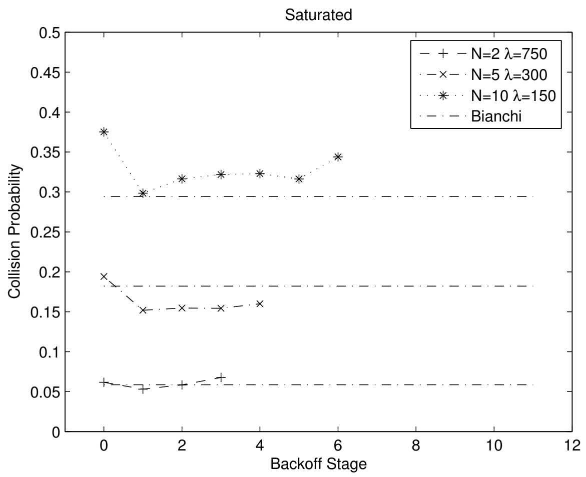

Figure: Saturated collision probabilities. Experimental data.

Testing (A2): $\{C_k\}$ identically distributed

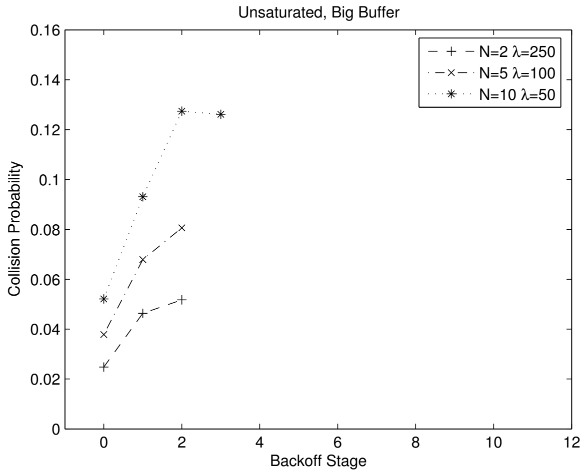

Figure: Unsaturated, big buffer collision probabilities. Experimental data.

What are the big-buffer hypotheses?

Big-buffer models:

- $Q_k = 1$ if packet waiting after k^{th} successful transmission.

What are the big-buffer hypotheses?

Big-buffer models:

- $Q_k = 1$ if packet waiting after k^{th} successful transmission.
- $Q_k = 0$ if no packet waiting after k^{th} successful transmission.

What are the big-buffer hypotheses?

Big-buffer models:

- $Q_k = 1$ if packet waiting after k^{th} successful transmission.
- $Q_k = 0$ if no packet waiting after k^{th} successful transmission.

Assumptions:

- (A3) $\{Q_k\}$ is an independent sequence;

What are the big-buffer hypotheses?

Big-buffer models:

- $Q_k = 1$ if packet waiting after k^{th} successful transmission.
- $Q_k = 0$ if no packet waiting after k^{th} successful transmission.

Assumptions:

- (A3) $\{Q_k\}$ is an independent sequence;
- (A4) $\{Q_k\}$ are identically distributed with $P(Q_k = 1) = q$.

Testing (A3): $\{Q_k\}$ independent

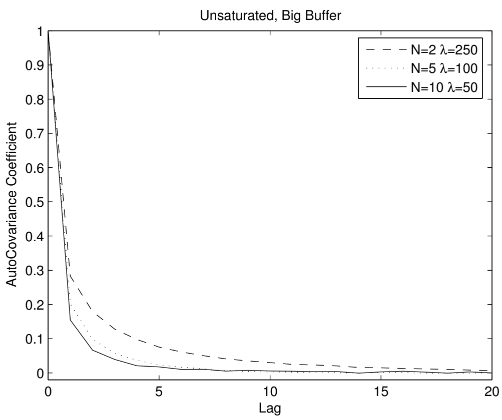

Figure: Unsaturated, big buffer queue-non-empty sequence normalized auto-covariances. Experimental data. $K = 1700k, 720k, 360k$.

Testing (A4): $\{Q_k\}$ identically distributed

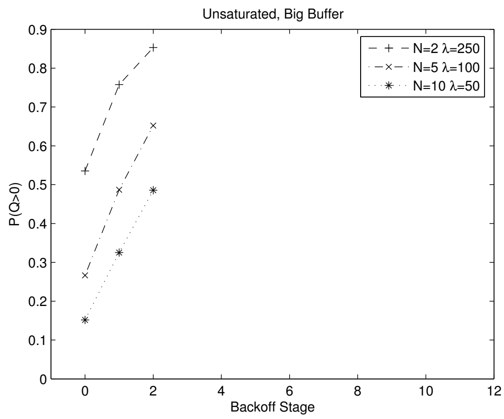

Figure: Unsaturated, big buffer queue-non-empty probabilities. Experimental data. (Note the large y-range!)

What are the 802.11e hypotheses?

Models with different AIFS values:

- H_k is length of k^{th} stuck in a hold-state.

What are the 802.11e hypotheses?

Models with different AIFS values:
- H_k is length of k^{th} stuck in a hold-state.

Assumptions:
- (A5) $\{H_k\}$ is an independent sequence;

What are the 802.11e hypotheses?

Models with different AIFS values:
- H_k is length of k^{th} stuck in a hold-state.

Assumptions:
- (A5) $\{H_k\}$ is an independent sequence;
- (A6) $\{H_k\}$ are identically distributed with a distribution that can be determined from a stopping time problem.

Testing (A5): $\{H_k\}$ independent

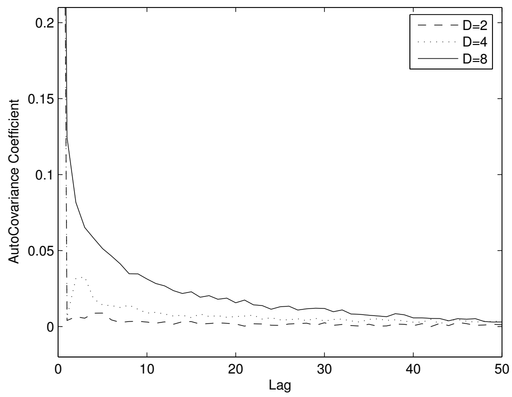

Figure: Hold state normalized auto-covariances. 5 class 1, 5 class 2 stations, $D = 2, 4 \& 8$. $K = 1700k, 1200k, 850k$. ns-2 data

Testing (A6): $\{H_k\}$ specific distribution

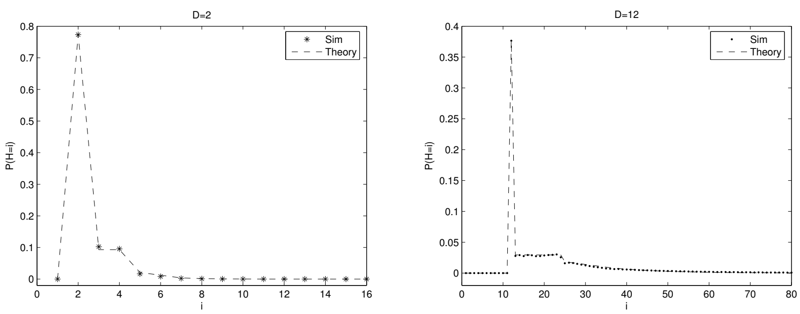

Figure: Hold state distributions, $D = 2, 12$. ns-2 data.

Testing (A6): $\{H_k\}$ specific distribution

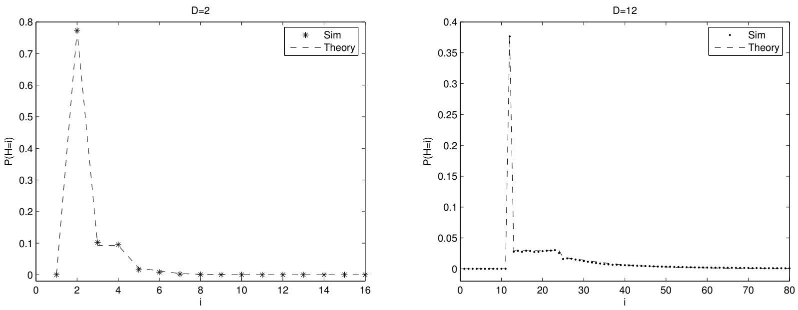

Figure: Hold state distributions, $D = 2, 12$. ns-2 data.

Kolmogorov-Smirnov test accepts fit for K of the order $10,000$;

Testing (A6): $\{H_k\}$ specific distribution

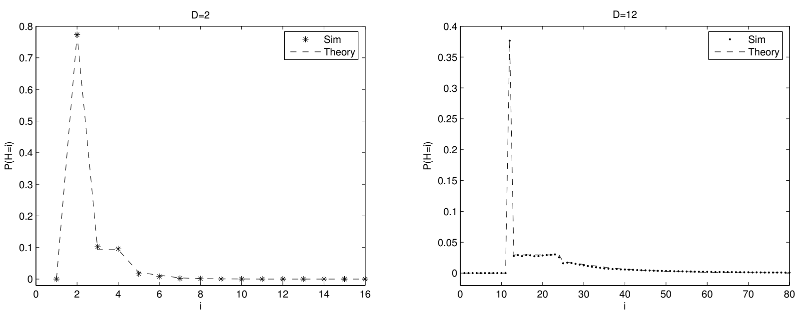

Figure: Hold state distributions, $D = 2, 12$. ns-2 data.

Kolmogorov-Smirnov test accepts fit for K of the order $10,000$;
rejects it for K of the order $1,000,000$.

What are the 802.11s hypotheses?

Mesh model(s) assume:

- D_k is k^{th} inter-departure time.

What are the 802.11s hypotheses?

Mesh model(s) assume:

- D_k is k^{th} inter-departure time.

Assumptions:

- (A7) $\{D_k\}$ is an independent sequence;

What are the 802.11s hypotheses?

Mesh model(s) assume:

- D_k is k^{th} inter-departure time.

Assumptions:

- (A7) $\{D_k\}$ is an independent sequence;
- (A8) $\{D_k\}$ are exponentially distributed.

Summary

Assumption	Sat.	Small buf.	Big buf.
(A1) $\{C_k\}$ indep.	✓	✓	✓
(A2) $\{C_k\}$ i. dist.	✓	✓	✓/×
(A3) $\{Q_k\}$ indep.	-	-	✓/×
(A4) $\{Q_k\}$ i. dist.	-	-	×
(A5) $\{H_k\}$ indep.	✓/×	-	-
(A6) $\{H_k\}$ dist.	✓	-	-
(A7) $\{D_k\}$ indep.	✓	✓	✓
(A8) $\{D_k\}$ exp. dist.	×	✓	✓

Table: $\{C_k\}$ collision sequence; $\{Q_k\}$ queue-occupied sequence; $\{H_k\}$ hold sequence; $\{D_k\}$ inter-departure time sequence.

Summary

Assumption	Sat.	Small buf.	Big buf.
(A1) $\{C_k\}$ indep.	✓	✓	✓
(A2) $\{C_k\}$ i. dist.	✓	✓	✓/×
(A3) $\{Q_k\}$ indep.	-	-	✓/×
(A4) $\{Q_k\}$ i. dist.	-	-	×
(A5) $\{H_k\}$ indep.	✓/×	-	-
(A6) $\{H_k\}$ dist.	✓	-	-
(A7) $\{D_k\}$ indep.	✓	✓	✓
(A8) $\{D_k\}$ exp. dist.	×	✓	✓

Table: $\{C_k\}$ collision sequence; $\{Q_k\}$ queue-occupied sequence; $\{H_k\}$ hold sequence; $\{D_k\}$ inter-departure time sequence.

K. D. Huang, K.D & D. Malone, Tech. Report.
(Preliminary report: K. D. Huang, K.D, D. Malone & D. Leith, IEEE PIMRC 2008.)

Epilogue: Impact of erroneous hypotheses?

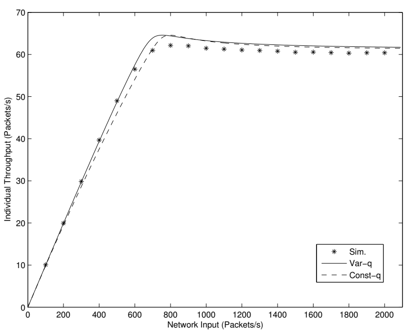

Epilogue: Impact of erroneous hypotheses?

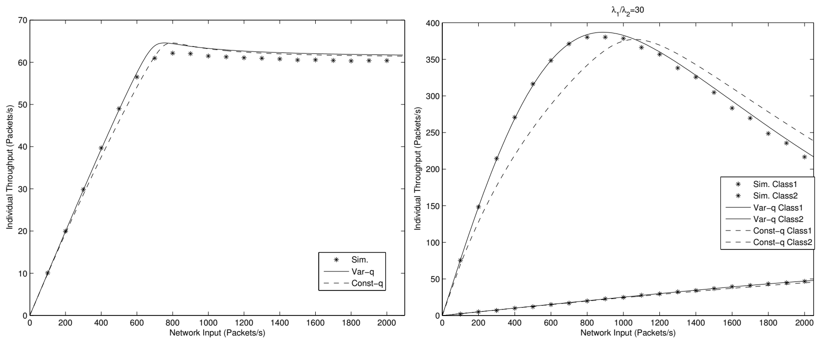

Figure: Theory & ns-2 data.

K. D. Huang & K.D, IEEE Comms Letters 2009.

Conclusions

Assumption	Sat.	Small buf.	Big buf.
(A1) $\{C_k\}$ indep.	✓	✓	✓
(A2) $\{C_k\}$ i. dist.	✓	✓	×
(A3) $\{Q_k\}$ indep.	-	-	×
(A4) $\{Q_k\}$ i. dist.	-	-	×
(A5) $\{H_k\}$ indep.	✓/×	-	-
(A6) $\{H_k\}$ dist.	✓	-	-
(A7) $\{D_k\}$ indep.	✓	✓	✓
(A8) $\{D_k\}$ exp. dist.	×	✓	✓

Conclusions

Assumption	Sat.	Small buf.	Big buf.
(A1) $\{C_k\}$ indep.	✓	✓	✓
(A2) $\{C_k\}$ i. dist.	✓	✓	×
(A3) $\{Q_k\}$ indep.	-	-	×
(A4) $\{Q_k\}$ i. dist.	-	-	×
(A5) $\{H_k\}$ indep.	✓/×	-	-
(A6) $\{H_k\}$ dist.	✓	-	-
(A7) $\{D_k\}$ indep.	✓	✓	✓
(A8) $\{D_k\}$ exp. dist.	×	✓	✓

Reports available at:

`http://www.hamilton.ie/ken_duffy`