Energy-Conserving Privacy-Enhancing Algorithms in Resource-Constrained Devices

Michael M. Groat
Department of Computer Science
University of New Mexico

July 27, 2012

Chair:	Dr. Stephanie Forrest
Co-Chair:	Dr. Wenbo He
Committee:	Dr. Fernando Esponda
	Dr. Terran Lane
	Dr. Jared Saia

What are Resource-Constrained Devices?

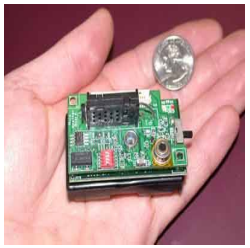

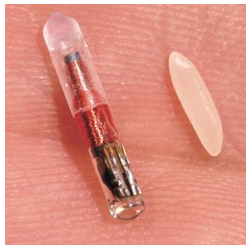

- Small, ubiquitous sensors that are common in our physical environment.
  - Limitations of memory, energy, radio range, CPU complexity, etc.
- Or, larger personal devices that have a limited battery, e.g., smart phones.

- These devices increasingly collect human related data, which needs to be secured.

- The standard methodology is encryption.

Problem: Encryption is Expensive

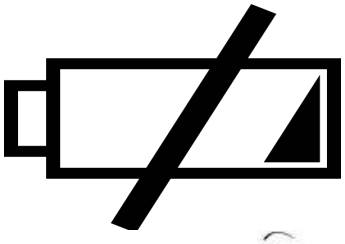

1) Energy expensive [1].

2) Slow.

3) Must ultimately trust a final recipient.

This dissertation addresses these problems, using non-cryptographic methods.

[1] C. Karlof, et al., “Tinysec: a link layer security architecture for wireless sensor networks,” in SenSys ’04, pp. 162–175.

Why is this Problem Important?

- People's privacy matters.
- Energy consumption.

- Protecting the privacy of unwilling participants caught in the growth of ubiquitous sensors is also important.

Key Idea: Data Indistinguishability

- **Definition:** A single element in a set of size k is k-indistinguishable if cannot be identified with a higher probability than random guessing.
- I apply indistinguishability in novel ways to resource-constrained devices.

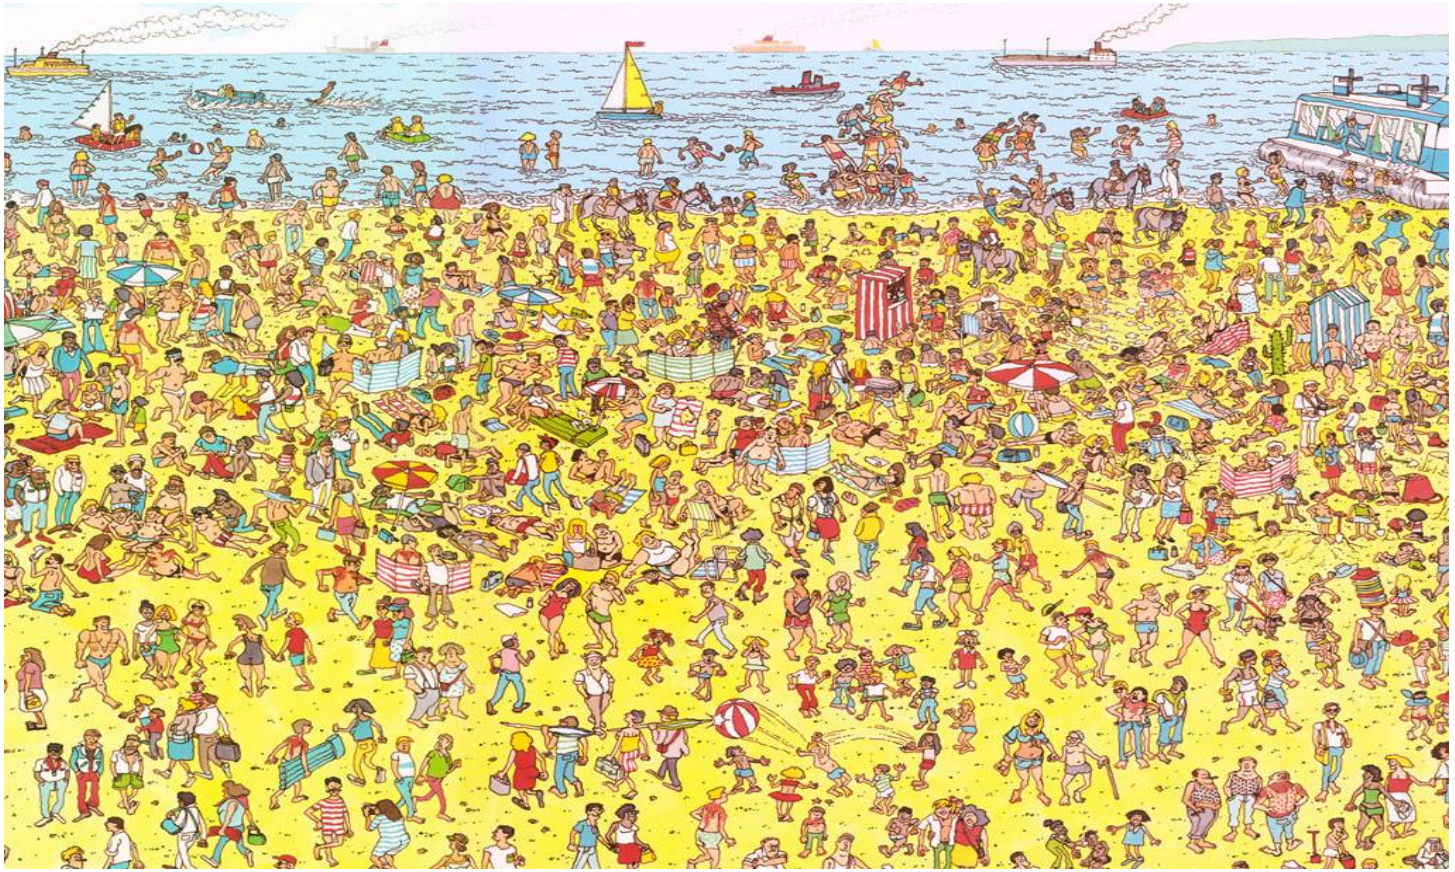

Images from www.findwaldo.com

Dissertation Overview

- Thesis statement: If the privacy requirement can be relaxed to indistinguishability then:
  - Sensor data can be protected.

    (and)
  - Energy can be saved.

- I achieve this goal with two information collection protocols:

  1) Multidimensional negative surveys (MDNSs).

  2) k-Indistinguishable privacy preserving data aggregation (KIPDA).

- The protocols are studied with analysis, simulations, and implementations.

Chapter 2: Background & Preliminary Work

Negative Surveys

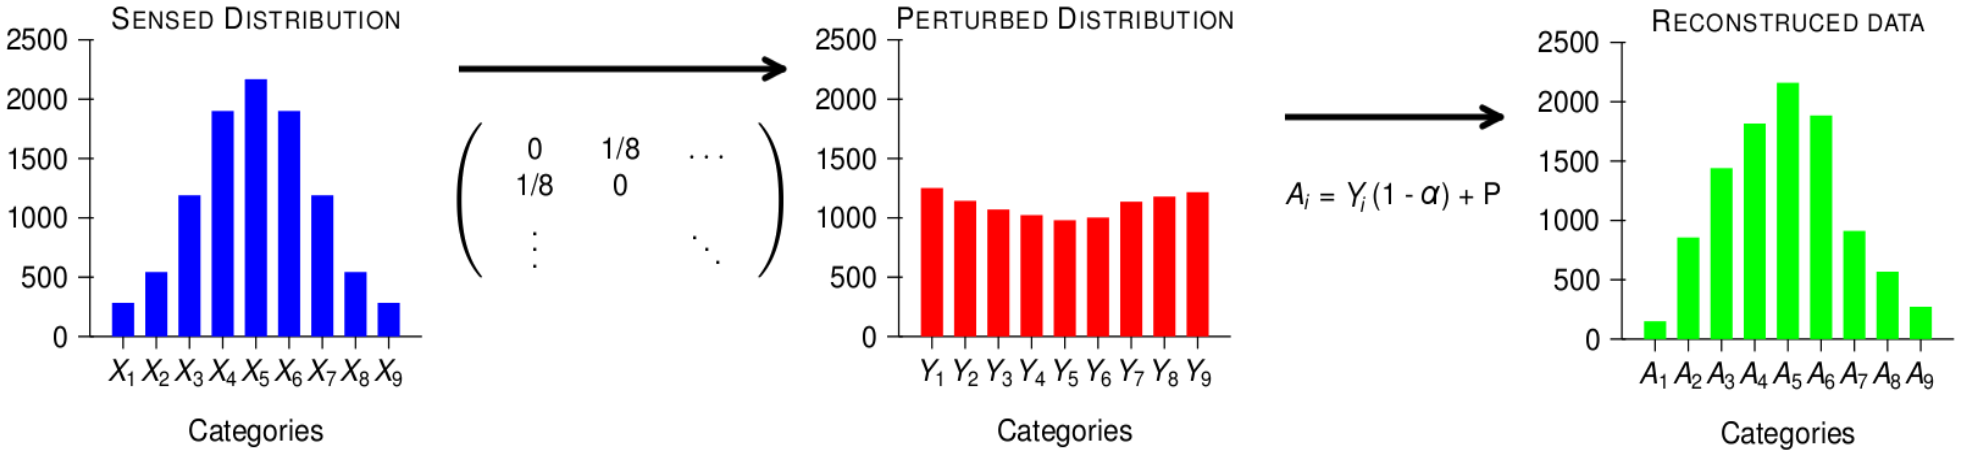

Positive Survey:
Which disease do you have?

- [x] X_1, Diabetes
- [ ] X_2, Cancer
- [ ] X_3, Heart disease
- [ ] X_4, Mood disorder

Negative Survey:
Which disease do you *not* have?

- [ ] Y_1, Diabetes
- [ ] Y_2, Cancer
- [x] Y_3, Heart disease
- [ ] Y_4, Mood disorder

- Preliminary work done in collaboration with James Horey.
- An extension of random response techniques from privacy preserving data mining:
  1) Each sensed sample (of 10,000 samples) is perturbed according to a matrix of probabilities.
  2) The original distributions is estimated using reconstruction algorithms.

Equation for the one-dimensional case is efficient. However, in the multidimensional case it is exponential with the number of dimensions.

Privacy-preserving Data Aggregation (PDA)

- Wireless sensor network (WSN)

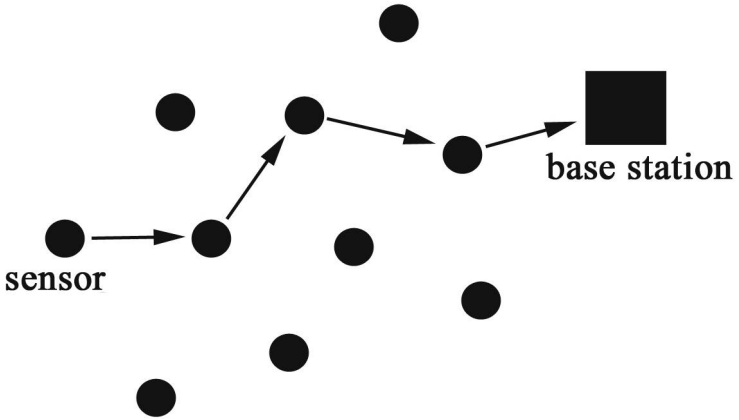

Data collection:

No data aggregation

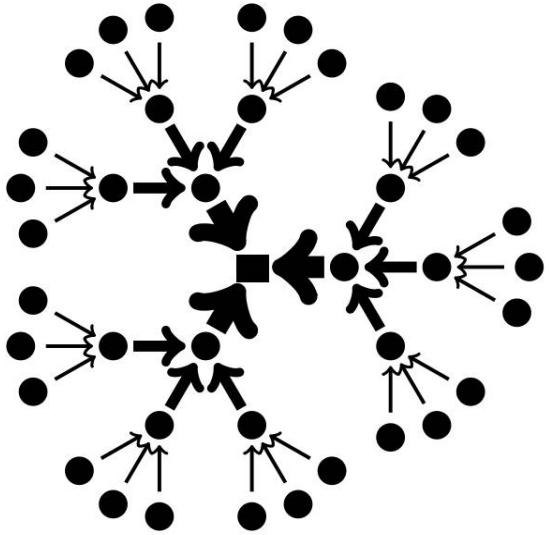

With data aggregation

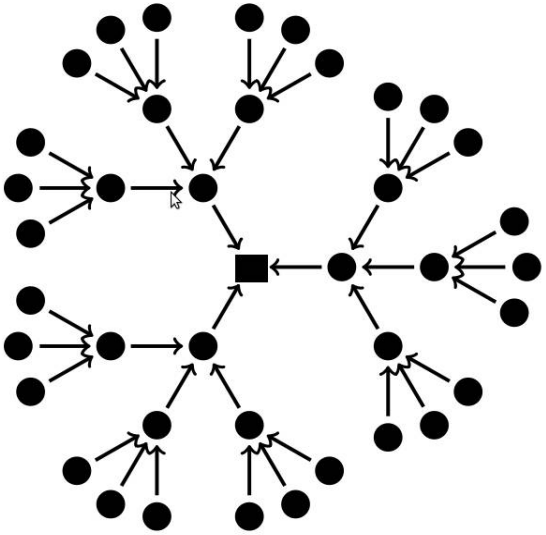

- Data aggregation is trivial, until privacy adds the following challenges:
  1) Untrusted nodes in the network.
  2) Untrusted base station(s).
  3) Non-linear aggregation functions.
  4) Energy budget.

- Two main-stream solutions: hop-by-hop and end-to-end encryption.

Privacy Assumptions and Definitions

- Privacy notions:
  - Indistinguishability (KIPDA).
  - Perturbation (MDNSs).
  - Not binary.

- Threat model:
  - Level 1: Eavesdroppers intercepting packets.
  - Level 2: Honest but curious [2] nodes in the network.
  - Level 3: Honest but curious base station.

- Assumptions:
  - Only a small number of nodes or data are compromised.
  - Polynomial time adversaries based on their input size.

[2] O. Goldreich, Foundations of Cryptography: Volume 2, Basic Applications. New York, NY: Cambridge University Press, 2004.

Chapter 3: Multidimensional Negative Surveys (MDNSs)

MDNS Protocols

	a	b	c
1			
2			
3			
4		X	
5			

- Node protocol:
  1) The sensor records data.
  2) For each data dimension, the sensor chooses a category other than what was sensed.
  3) Transmits the negative data.

- Base station protocol:
  1) (Optional) Request the data.
  2) Collect the data.
  3) Estimate the sensed data from the negative values with a reconstruction algorithm.

Reconstruction Algorithm

- ***Problem:*** Natural extension from the one-dimensional case has exponential time with respect to the number of dimensions.
  - Example with two dimensions:

$$A(i,j)=\sum_{a=1,b=1}^{\alpha_1,\alpha_2} R(a,b)-(\alpha_1-1)\cdot\sum_{a=1}^{\alpha_1} R(a,j)-(\alpha_2-1)\cdot\sum_{b=1}^{\alpha_2} R(i,b)$$
$$+(\alpha_i-1)\cdot(\alpha_j-1)\cdot R(i,j)\mid \forall i,j$$

- ***Solution:*** Matrix memoization.
  - Partial results are cached back into the perturbed data to be reused.

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

Perturbed data in two dimensions, Y(:,:)

957	982	1010
632	664	731
719	718	646
1010	991	938

A matrix multiplication with a vector is performed:

$$R(:,1)=\begin{bmatrix} 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \\ 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \end{bmatrix}^{-1} \cdot \begin{bmatrix} 957 \\ 632 \\ 719 \\ 1010 \end{bmatrix}$$

$$O(\alpha^2)$$

But since this is a negative survey reconstruction, the single dimensional reconstruction equation is faster.

$$R(j,1)=P+(1-\alpha)\cdot R_{j,1} \mid \forall\, j$$

$$O(\alpha)$$

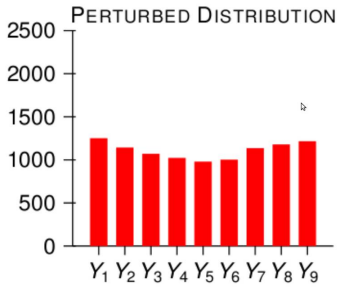

Y(:,:) is similar to the perturbed one-dimensional histogram seen earlier.

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

447	982	1010
1422	664	731
1161	718	646
288	991	938

R = perturbed data

And the results are stored back in R.

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

447	405	1010
1422	1365	731
1161	1203	646
288	384	938

R = perturbed data

We continue in the same dimension

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

447	405	295
1422	1365	1132
1161	1203	1387
288	384	511

R = perturbed data

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

253	337	557
1422	1365	1132
1161	1203	1387
288	384	511

R = perturbed data

The stored information is then used in the other dimensions.

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

253	337	557
1075	1189	1655
1161	1203	1387
288	384	511

R = perturbed data

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

253	337	557
1075	1189	1655
1429	1345	977
288	384	511

R = perturbed data

Reconstruction Algorithm

- Example: 2 dimensions of 4 and 3 categories, and 10,000 samples.

253	337	557
1075	1189	1655
1429	1345	977
607	415	161

R = perturbed data

Reconstruction Algorithm

RMSE = 81.45. 10,000 samples.

253	337	557
1075	1189	1655
1429	1345	977
607	415	161

Reconstructed data

202	391	664
988	1287	1468
1468	1287	988
664	391	202

Original sensed distribution

Privacy and Utility Metrics

- Originally used the relative root mean square error, now I use the following:

- Privacy Metric:
  - Measures the probability of guessing the original data from the disguised values, and is based on the maximum a posteriori estimate.

- Utility Metric:
  - Measures the difference between the original and reconstructed data distributions with the mean square error.

- For both metrics a lower value is more desirable. Privacy ranges from [0,1], while utility ranges from $[0, +\infty)$.

Dimensional Adjustment

- Addresses problem of having too few participants to maintain a given level of utility.
- Sacrifices a smaller amount of privacy for a greater gain in utility.

	1	2	3
1	a	b	X
2	d	e	f
3	g	h	i

a	
b	
c	X
d	
e	
f	
g	
h	
i	

- 1 dimension of 9 categories can be adjusted to 2 dimensions of 3 categories each.
- Trade-off example:

	1 dim 10,000 cats.	6 dims. 5x5x5x5x4x4 cats.
Utility	0.00100	0.00014
Privacy	0.01457	0.01960

- Privacy degrades 34% while utility increases 86%.

Radiation Detection Simulation

- Assume cell phones equipped with radiation detectors.
- Phones perturb their location and radiation level.
- Then report this information to a base station

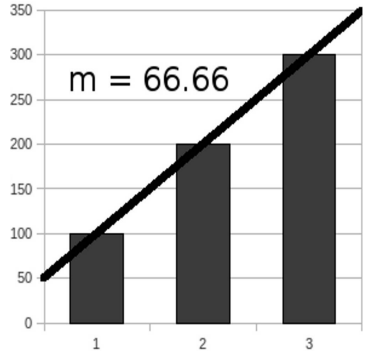

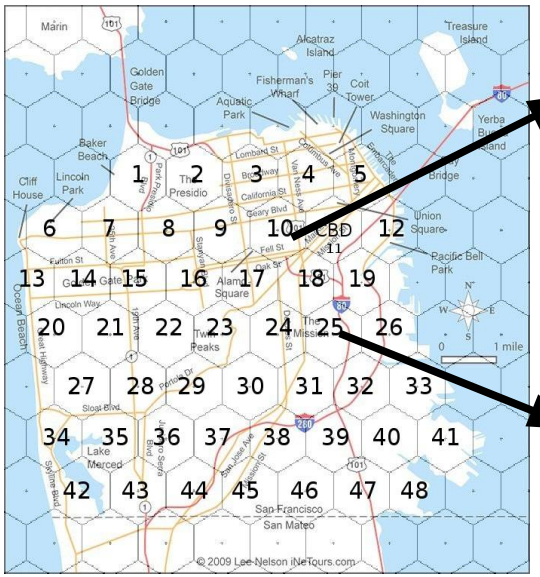

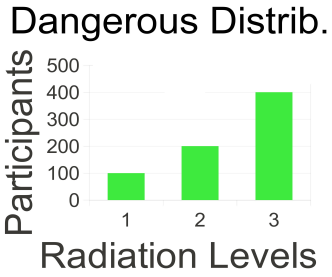

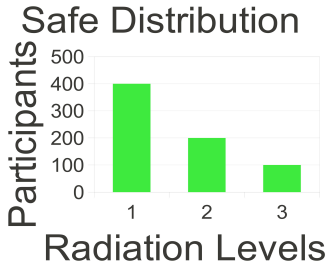

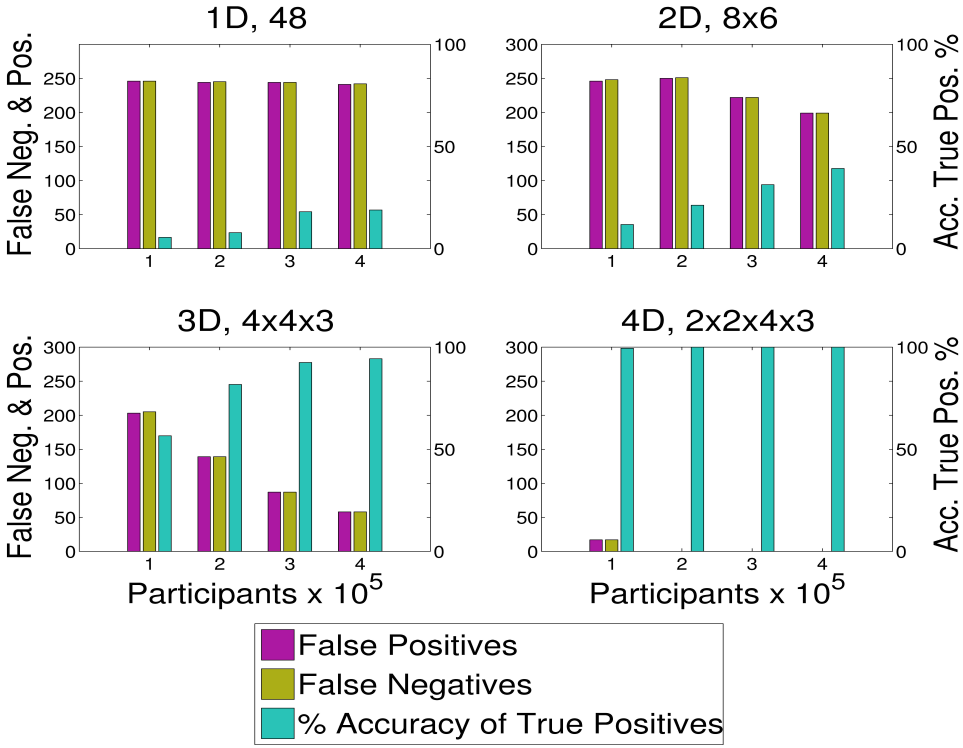

MDNSs on Android Smart Phones

- Physical cell phones were programmed with the MDNS node protocol, and a server with the base station protocol.
- Categories included:
  - 4 latitudes
  - 6 longitudes, dimensionally adjusted to 3 and 2 categories
  - 3 sound levels, sampled from the phone's microphone.

Original sensed distribution

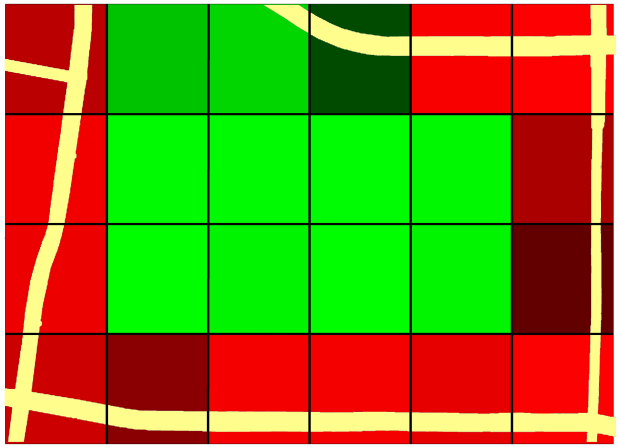

Reconstructed distribution

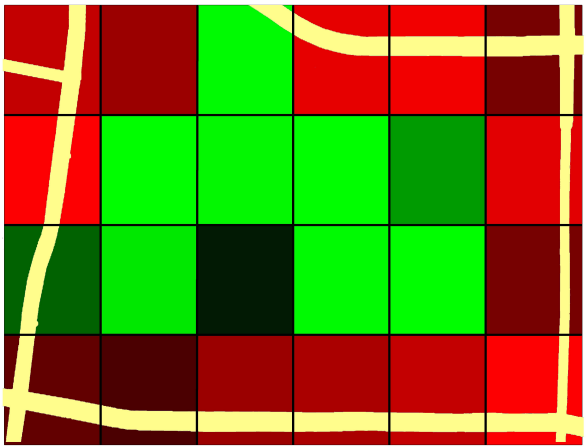

- Energy is conserved because encryption and key management and distribution are eliminated.
- 7300 samples with 72 overall categories. RMSE is 86.43.

MDNS: Impact of Research

- Xie et al.[1] propose a special perturbation matrix for Gaussian negative surveys.
- Quercia et al.[2] propose a yes/no RRT for every location.

[1] Xie et al. “Privacy-aware Collection of Aggregate Spatial Data”. Data Knowledge Engineering, 2011, Vol. 70, No. 6, pp. 576-595.

[2] Quercia et al. “SpotME If You Can: Randomized Responses for Location Obfuscation on Mobile Phones”, ICDCS, 2011, pp 363-372.

Chapter 4: KIPDA

PDA: Literature Review

	MAX/MIN	Accurate	Efficient	Level 1 Privacy	Level 2 Privacy
Girao et al.	No	Yes	Yes	Yes	Yes
Acharya et al.	Yes	Yes	Yes	Yes	No
Ertual et al.	Yes	Yes	No	Yes	Yes
Zhang et al.	Yes	No	Yes	Yes	Yes
KIPDA	Yes	Yes	Yes	Indisting.	*k*

- Girao et al. [3], and other mainstream end-to-end encryption schemes only support addition and/or multiplication aggregation functions.
- Rivest [4] showed *homomorphic privacy (HP)* is insecure against ciphertext only attacks, if a comparison operator, such as the MAX or MIN functions, is supported.
- Acharya et al. [5] use OPES (a type of HP) in WSNs. However, it is insecure for Level 2 privacy.
- Zhang et al. [6] uses histograms with hashed message authentication code (HMAC) messages. However, it is secure but not accurate.
- Ertual and Vaidehi [7] use additive HP to calculate comparison operations. However, it is accurate but not efficient.

KIPDA: Introduction

- I propose an energy-efficient k-indistinguishable MAX/MIN aggregation scheme that satisfies:
  - Level 1 privacy with indistinguishability.
  - Level 2 privacy that scales.

- Aggregates are hidden in plain sight among camouflage data inside a message set (vector).

- The positions in the message set obey special properties and vary between nodes, such as dividing them into restricted and unrestricted sets.

KIPDA: Example

- Example of aggregation:

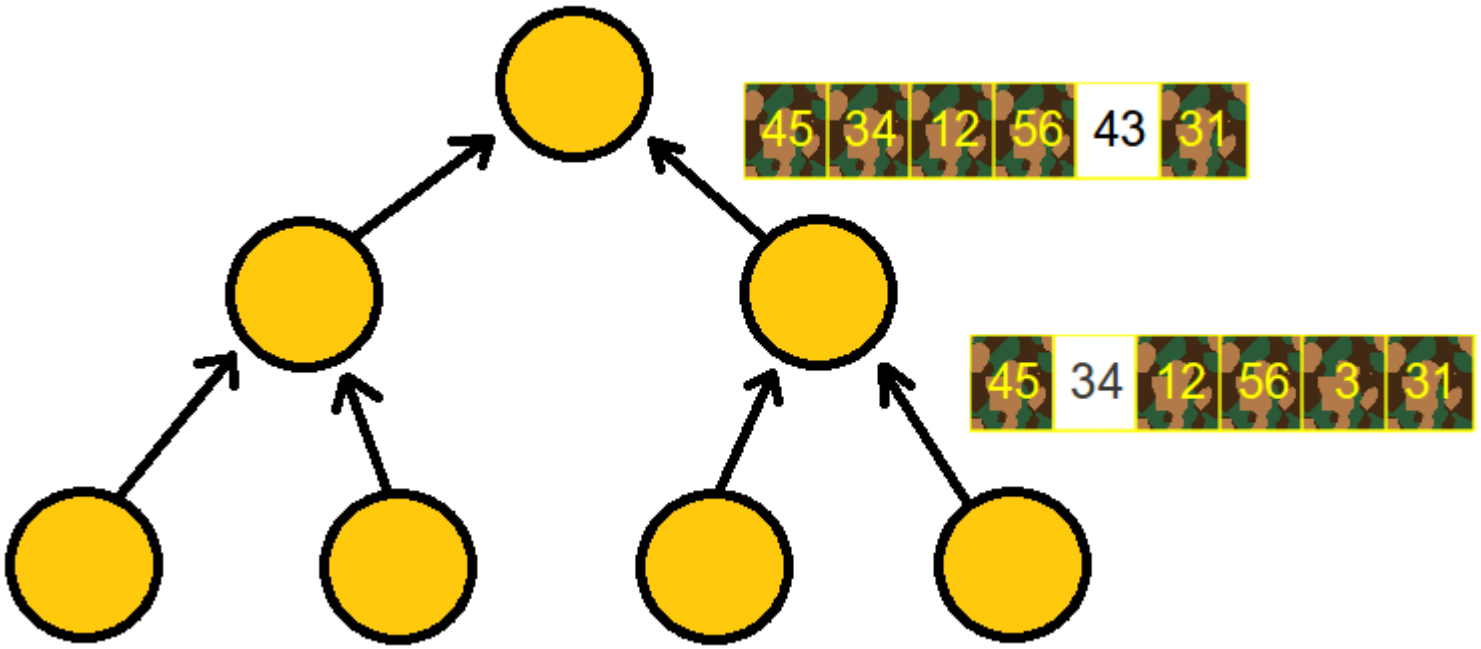

KIPDA: Properties

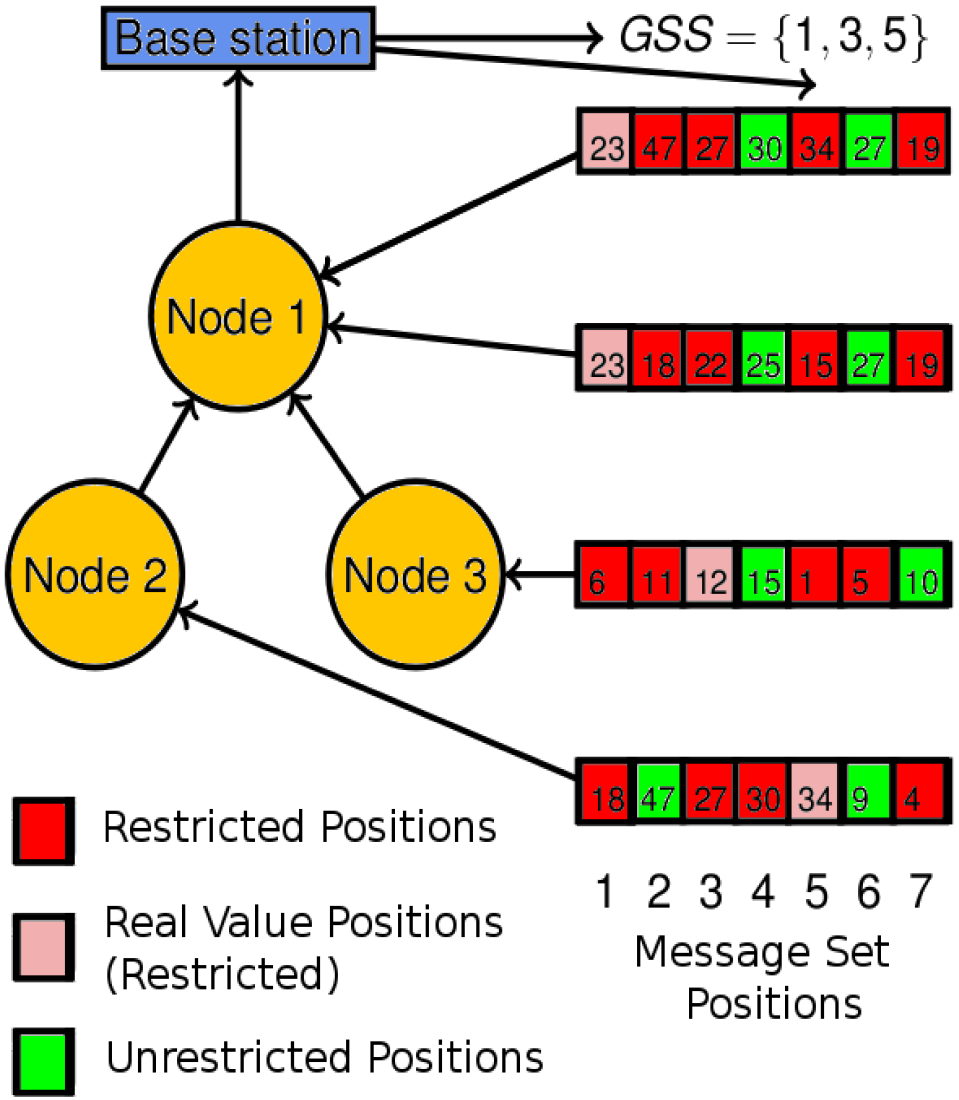

- Property 1:
  The Real Value Position is a subset of the Global Secret Set (GSS).
- Property 2:
  The Restricted Value Positions must contain elements from both GSS and $\overline{GSS}$.
- Property 3:
  The Restricted Value Positions contain every element of GSS.
- Property 4:
  In the reporting phase, The Restricted Values must be less than or greater (MAX or MIN Agg.) than the Real Value. Unrestricted positions can be either.

KIPDA: Protocols

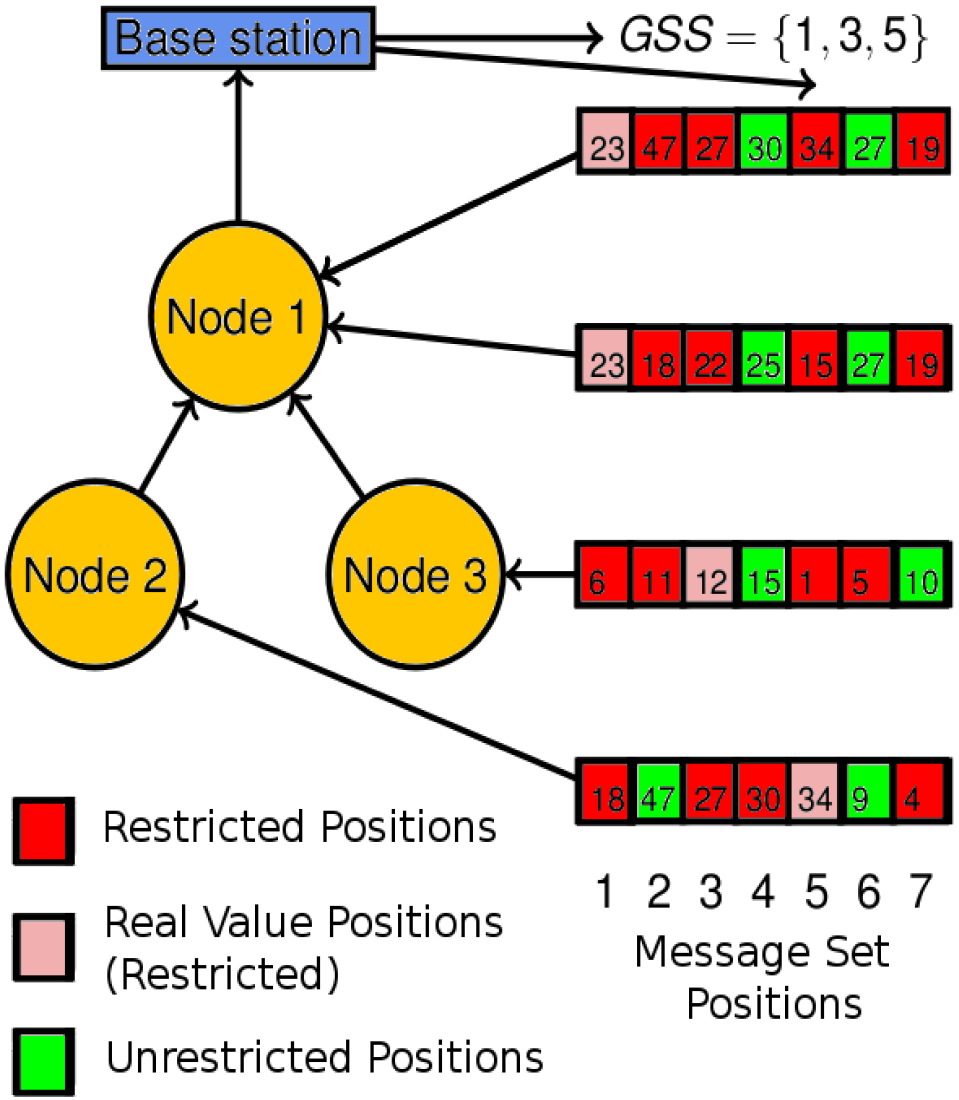

1) Pre-distribution

2) Reporting

3) Aggregation

4) Base station post processing

Optimal Sizing of Sets: GSS, NSS

Set sizes are determined in the following order:

1) The message set
   - According to energy available.
2) The node secret set *NSS*
   - According to protect against one rogue node or many nodes.
3) The global secret set *GSS*
   - According to the intersection of two lines that are based on the coupon collection problem.
     - The ability to pick GSS elements.
     - The ability to pick $\overline{\text{GSS}}$ elements.

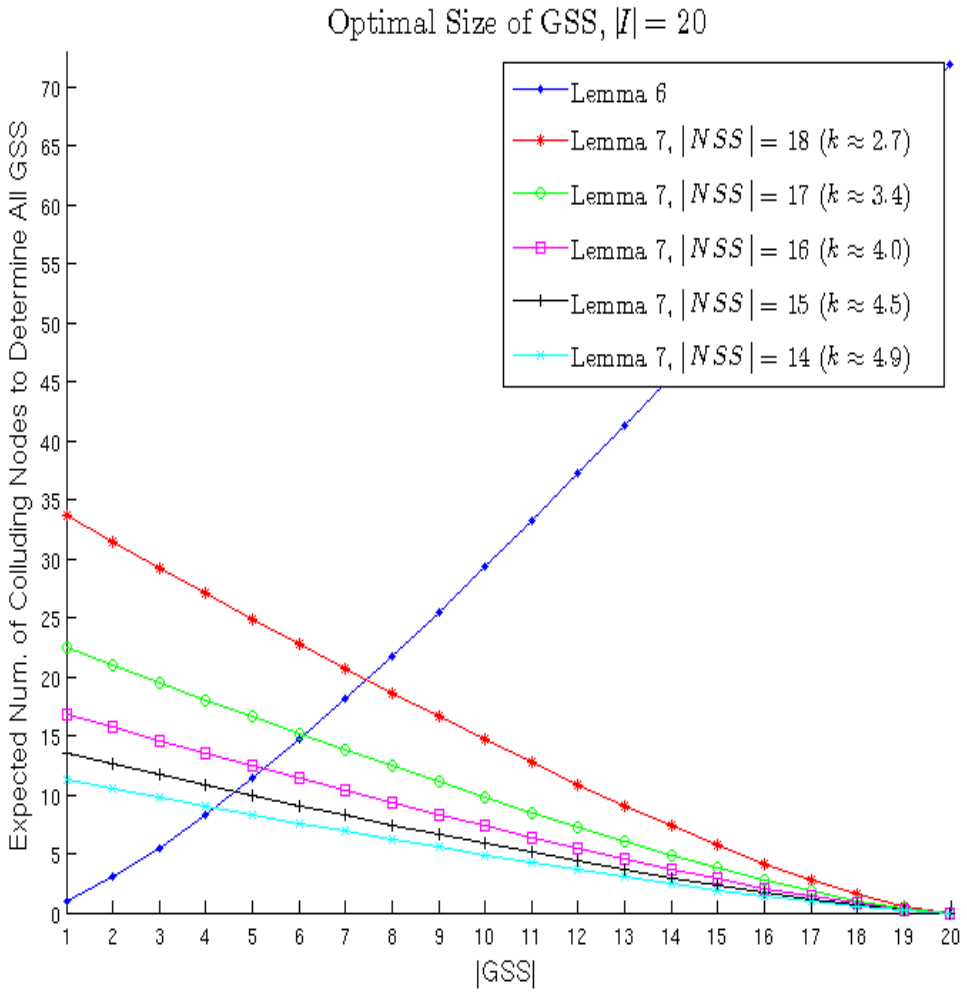

Determining *k*

- An upper bound on *k is*: $k \leq min(|\overline{NSS}|+1, |NSS|-1)$
- To an outside observer *k* is the size of the message set.
- I build, with a series of equations, the *k* value when *x* nodes collude (see figure).
- For space and time reasons, I have left these equations out. However, they are in the supplemental slides.

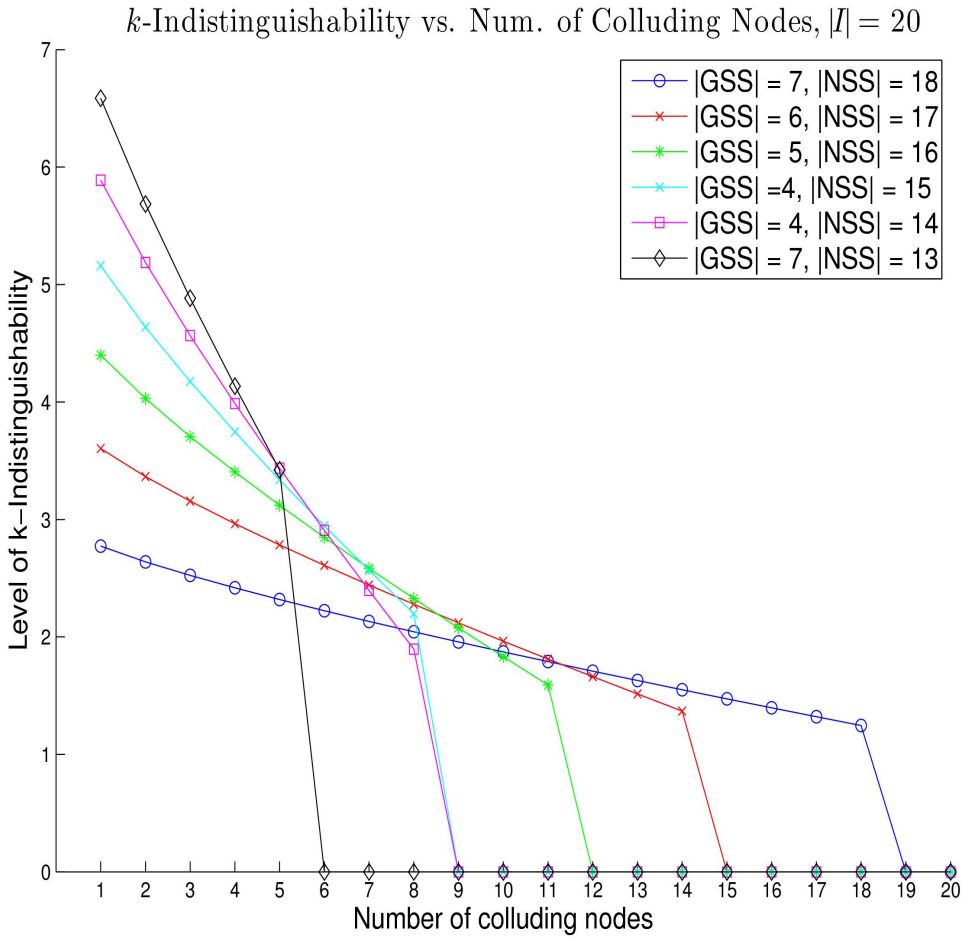

POWER-TOSSIM-Z Simulations

- Implemented KIPDA and hop-by-hop aggregation in TOSSIM, TinyOS simulator, with care to energy use.
- Compared energy use between KIPDA and hop-by-hop aggregation with AES, Skipjack, RC5, Trivium, and no encryption.

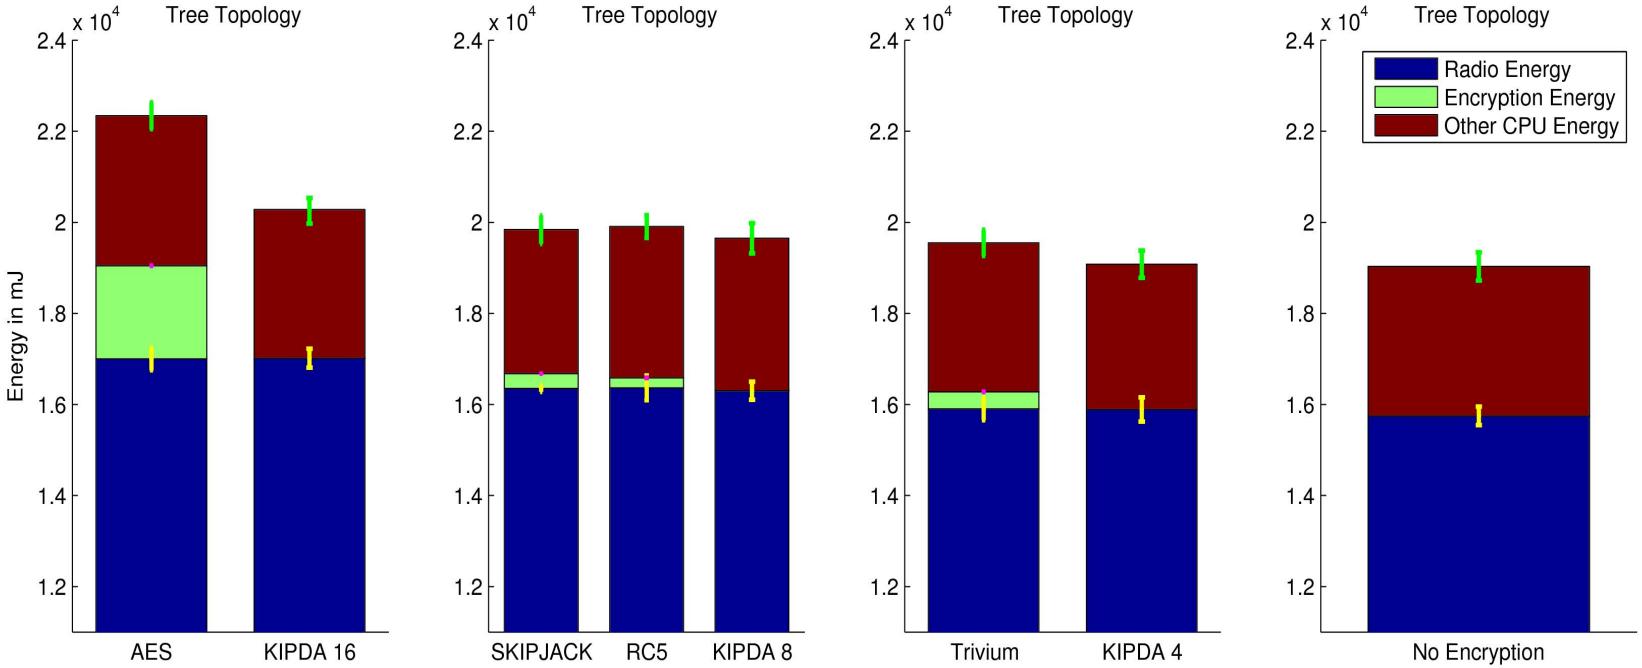

KIPDA Implementations

KIPDA was implemented on Moteiv T-mote Invent sensors.

- However, without physically altering the device, direct energy measurements were not possible.
- Voltage level could be obtained, but energy use cannot be completed.
- Ran three tests :
  - Voltage level between KIPDA and hop-by-hop aggregation with AES
  - Life time of each device in the network between KIPDA and hop-by-hop aggregation with AES.
  - Measured the time for each encryption primitive for AES, SkipJack, RC5, Trivium, and TinyECC

Discussion

- Other applications:
  - Pollution monitoring.
  - African malaria surveys.
- Trade-off analysis
  - MDNS: granularity, privacy, accuracy.
  - KIPDA: Privacy, efficiency, accuracy.
- Theoretical analysis overestimated encryption energy use.
  - However, I mostly used numbers and results from the literature.
  - TOSSIM simulations agree with other energy results in the literature using actual devices.

Limitations and Caveats

- MDNS
  - Rogue sensors that try to disrupt the aggregate.
  - Mobile sensors that don't move, or move between only a few locations.
  - MDNS do not perform well with more than 3,000 categories.

- KIPDA
  - Susceptible to statistical attacks if the sets GSS, NSS^i, and NSS^i_T are not regularly refreshed.
  - Network wide MIN and MAX needs to be kept secret.
  - Sensors work best in high entropy environments.

Future Work

- CDA
  - Secure knowledge of which node generated the MAX or MIN.
  - Privacy preserving fast aggregation (which KIPDA excels at).
- KIPDA
  - Variable sizes for the Restricted Value Positions and message set.
  - Compare energy use to other related work such as those in the related work matrix (Slide 29).
  - Nodes determine which are rogue.
- MDNS
  - Correlations between dimensions.
  - Other aggregates besides histograms
- General
  - Other areas where indistinguishability saves energy and preserves privacy.

Conclusion

- MDNSs
  - Protect location privacy
  - Still allow useful information to be collected.
- KIPDA:
  - Addresses MIN/MAX PDA efficiently and accurately.
  - Provides partial protection against in-network nodes.
  - Indistinguishability replaces encryption.
- Simulations and implementations show the feasibility, possibility, and tunability of the two presented protocols.

Publications

- In submission:
  - Michael M. Groat, Benjamin Edwards, James Horey, Wenbo He, Stephanie Forrest, **Applications and Analysis of Multidimensional Negative Surveys in Participatory Applications.** In submission to Pervasive and Mobile Computing Journal.
- Refereed Conferences:
  - Michael M. Groat, Benjamin Edwards, James Horey, Wenbo He, and Stephanie Forrest, **Enhancing Privacy in Participatory Sensing Applications with Multidimensional Data,** In Proceedings of the Tenth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom '12), March 2012, pp. 144-152, Lugano, Switzerland. (Acceptance ratio: 10.7% = 16/150)
  - Michael M. Groat, Wenbo He, and Stephanie Forrest, **KIPDA: k-Indistinguishable Privacy-preserving Data Aggregation in Wireless Sensor Networks**, In Proceedings of the Thirtieth Annual IEEE International Conference on Computer Communications (InfoCom '11), April 2011, pp. 2024-2032, Shanghai, China. (Acceptance ratio: 16.0% = 291/1823)
  - James Horey, Michael M. Groat, Stephanie Forrest, and Fernando Esponda, **Anonymous Data Collection in Sensor Networks**, In Proceedings of the Fourth Annual International Conference on Mobile and Ubiquitous Systems: Networking & Services (Mobiquitous '07), August 2007, pp. 1-8, Philadelphia, PA, USA. (Acceptance ratio: 22.7% = 27/119)
- Workshops:
  - James Horey, Stephanie Forrest, Michael Groat, **Reconstructing Spatial Distributions from Anonymized Locations**, In ICDE Workshop on Secure Data Management on Smartphones and Mobiles, April 2012. In Press, Washington DC, USA.

Thank you

Questions?

KIPDA Supplementary Material

Devices that Sense Low Entropy Distributions

- KIPDA:
  - Assume sensors only sense a high entropy environment.
  - Assume an adversary has no knowledge of the environment
  - Send a distribution in the message set that is tight around the sensed value.
  - Assume that the adversary knows just one node that senses a low entropy environment, and treat this node as rogue and only a partial amount of privacy is lost.
- MDNSs
  - Individuals sense low entropy distributions.
  - The sensor can keep track of negative information sent over time (see Information gained graph), and discontinue responses when a threshold has been reached. Then perhaps the individual can request a new unique user id from the base station?

Why Secure Multi-party Computation (SMC) is not Feasible

- MAX/MIN PDA and SMC are very similar
- SMC is a general case of Yao's millionaire problem.
- Solution to the problem is too resource intensive.
- Previous solutions [1,2] leverage public/private key cryptography.

[1] I. Damgard, M. Geisler, and M. Kroigard, “Homomorphic encryption and securecomparison,” International Journal of Applied Cryptograhpy, vol. 1, no. 1, pp. 22–31, 2008.

[2] A. C.-C. Yao, “Protocols for secure computations (extended abstract),” in 23rd Annual Symposium on Foundations of Computer Science (FOCS). IEEE, November 1982, pp. 160–164.

Comparison to Chaffing and Winnowing

- Chaffing and Winnowing concept.
- Requires 3rd party trust.
- Each bit requires a MAC, this is both expensive to create and expensive to transmit. (Especially per bit.)
- Camouflage is call chaff (like the wheat).
- Random number for chaff. Bob (Alice, Charles and Bob) knows the secret key shared with Alice, and know which MACs are valid.
- More for political reasons and export control.
- Method is not a homomorphic privacy technique, hence at best it could be used in hop-by-hop aggregation. However, I believe it is energy expensive.

KIPDA Analysis

- Results of the Analysis of Chapter 4.

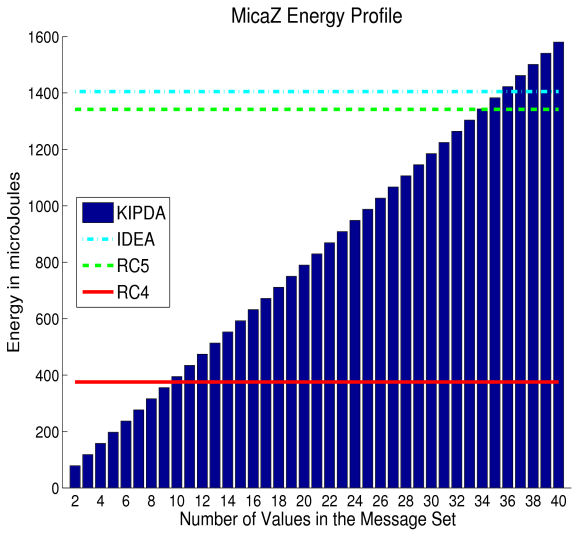

MicaZ

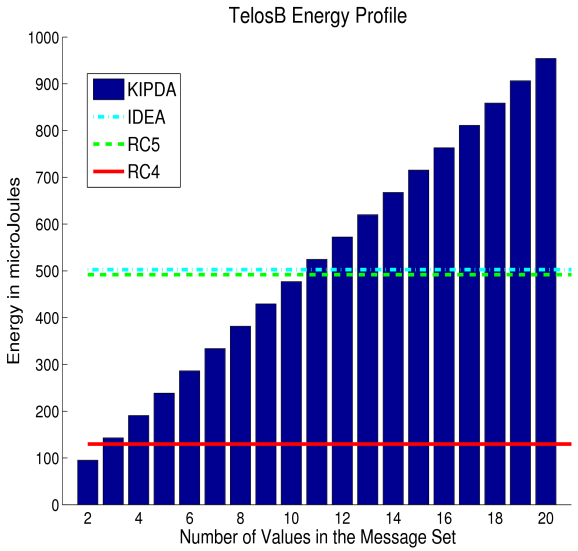

TelosB

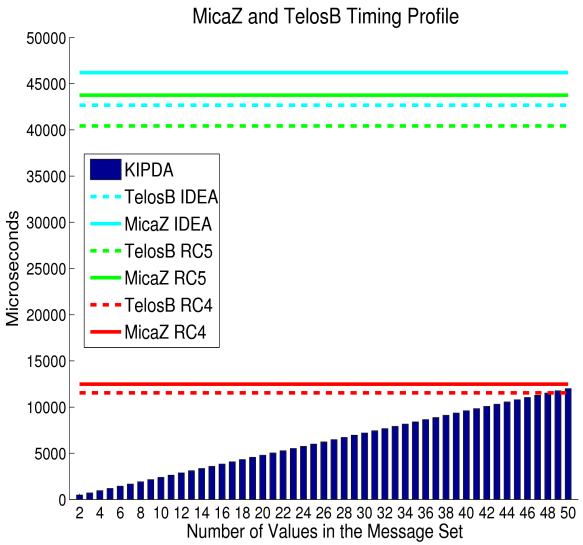

Timing - Both.

KIPDA Threshold Ratio

- Determines how large a message set can be sent that uses the same energy as hop-by-hop schemes.

$$\frac{E\left(Tx\left(m\,bits\right)+Rx\left(m\,bits\right)\right)}{E\left(Tx\left(l\,bits\right)+Rx\left(l\,bits\right)+Enc\left(l\,bits\right)+Dec\left(l\,bits\right)\right)}$$

- Where Tx is transmit, Rx is receive, and E is the energy of these events.
- Predicts network wide energy use.

KIPDA Simulation Design Steps

1) Query all nodes
2) All nodes respond to query
3) Aggregate the responses
4) Implement power readings and energy use.
5) Implement hop-by-hop encryption with the responses.
6) Implement KIPDA
7) Run simulations for various topologies, and types of encryption: AES, SkipJack, RC5, Trivium, and TinyECC

Average Value of *k*

- Average *k* value against a single rogue node.

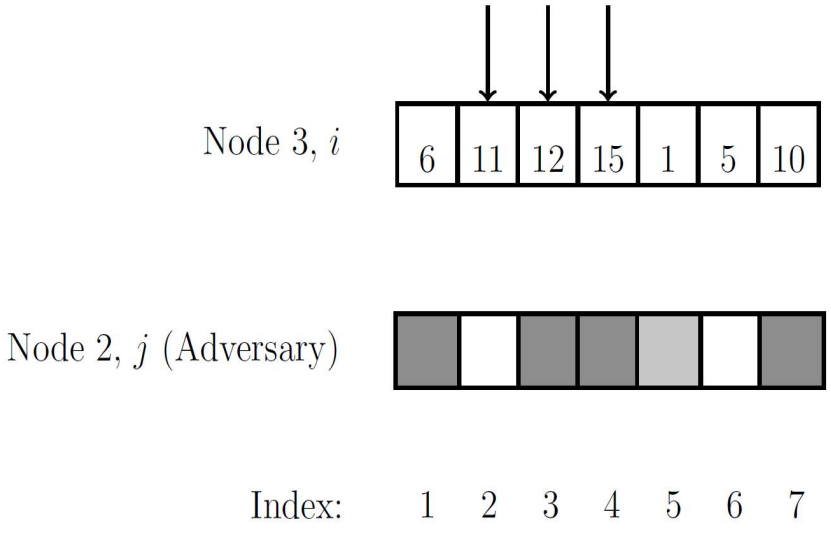

By $\overline{\text{NSS}}$

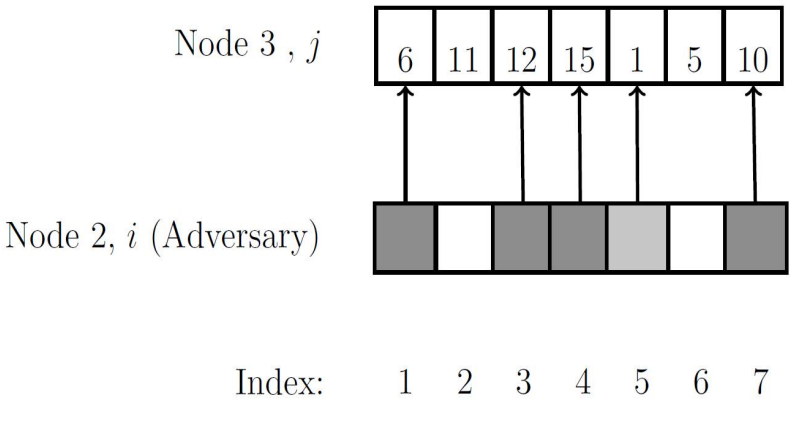

By NSS

k Values from *x* Colluding Nodes

$$k = \begin{cases} 0 & \text{if } |GSS_{known}| = |GSS| \text{ or } |\overline{GSS}_{Known}| = |\overline{GSS}| \\ min(\Lambda, \Pi) & \text{otherwise.} \end{cases}$$

$$\Pi = \sum_{g=1}^{|U|-|\overline{GSS}_{Known}|} P(g) \cdot (g-1)$$

$$P(g) = \frac{|GSS_{Known}|}{|U| - |\overline{GSS}_{known}|}.$$

$$\Lambda = \left(|\overline{NSS}| - E[\Psi] + 1\right),$$

$$E[\Psi] = \frac{|\overline{GSS}_{Known}| \cdot \left(|\overline{NSS}| + 1\right)}{|I|},$$

$$E[x] = |GSS| \cdot H_{|GSS|} = |GSS| \cdot \sum_{i=1}^{|GSS|} \frac{1}{i},$$

$$E[x] = \frac{|\overline{GSS}|}{|\overline{NSS}|} \cdot H_{|\overline{GSS}|} = \frac{|\overline{GSS}|}{|\overline{NSS}|} \cdot \sum_{i=1}^{|\overline{GSS}|} \frac{1}{i}.$$

$$|GSS_{known}| = \frac{x - \frac{1}{2}}{W(\frac{1}{2}e^{\gamma}(2x-1))}$$

$$|\overline{GSS}_{known}| = \frac{|\overline{NSS}| \cdot x - \frac{1}{2}}{W(\frac{1}{2}e^{\gamma}(2 \cdot |\overline{NSS}| \cdot x - 1))}.$$

Where γ (gamma) is the Euler-Mascheroni contant, 0.577215, and W is the Lambert W-Function, or product log.

MDNS Supplemental Material

Adversarial Information Gained

- 200 category SDNS, and a 20 by 10 2-dimensional negative survey. Max information to gain is $\log_2(200)$.

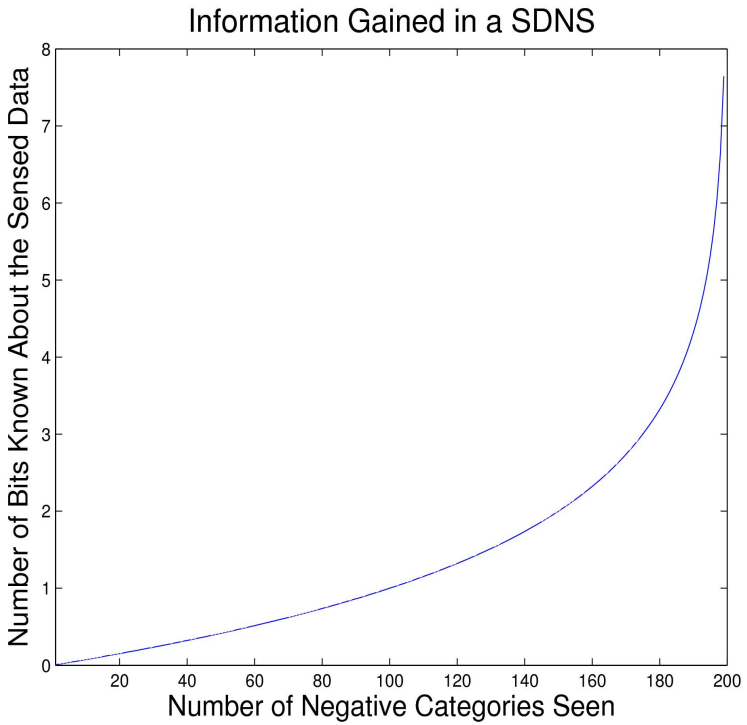

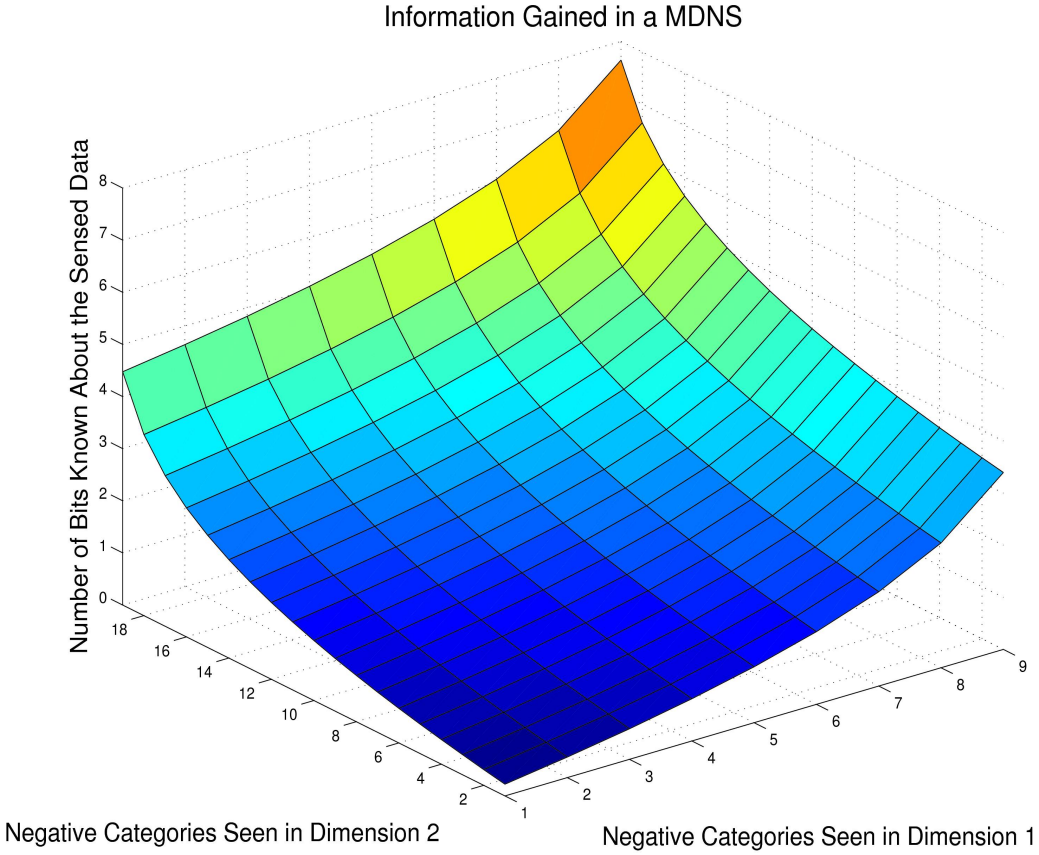

Comparison of Continuous MDNSs to Random Data Perturbation

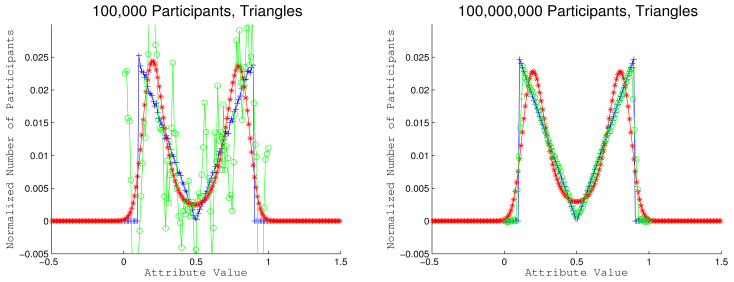

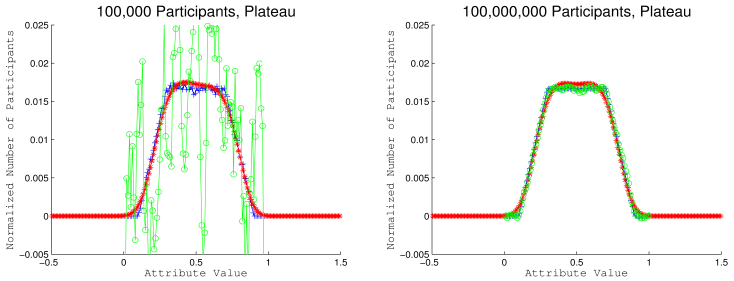

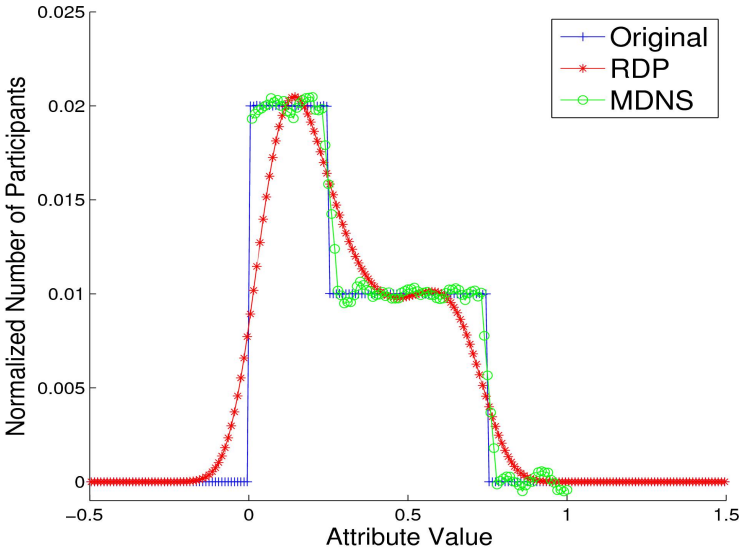

- Continuous negative surveys:
  - Original idea by Benjamin Edwards.
  - Each digit is a dimension of 10 categories, 0 to 9.
- Data from a distribution is perturbed and reconstructed at the base station for both RDP [1] and continuous negative surveys.
- Comparison to RDP shows strengths and weaknesses.

[1] R. Agrawal, and R. Skrikant. “Privacy-preserving Data Mining.” in Proc. of 2000 Int. Conf. of Management of Data. pp 439-450.

Effects of the Original Distribution

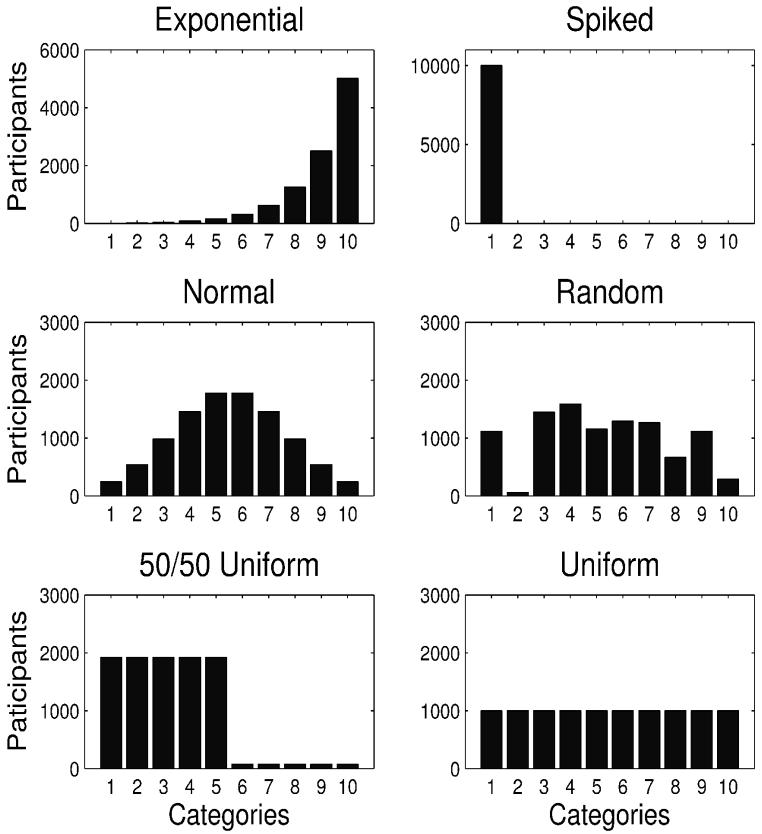

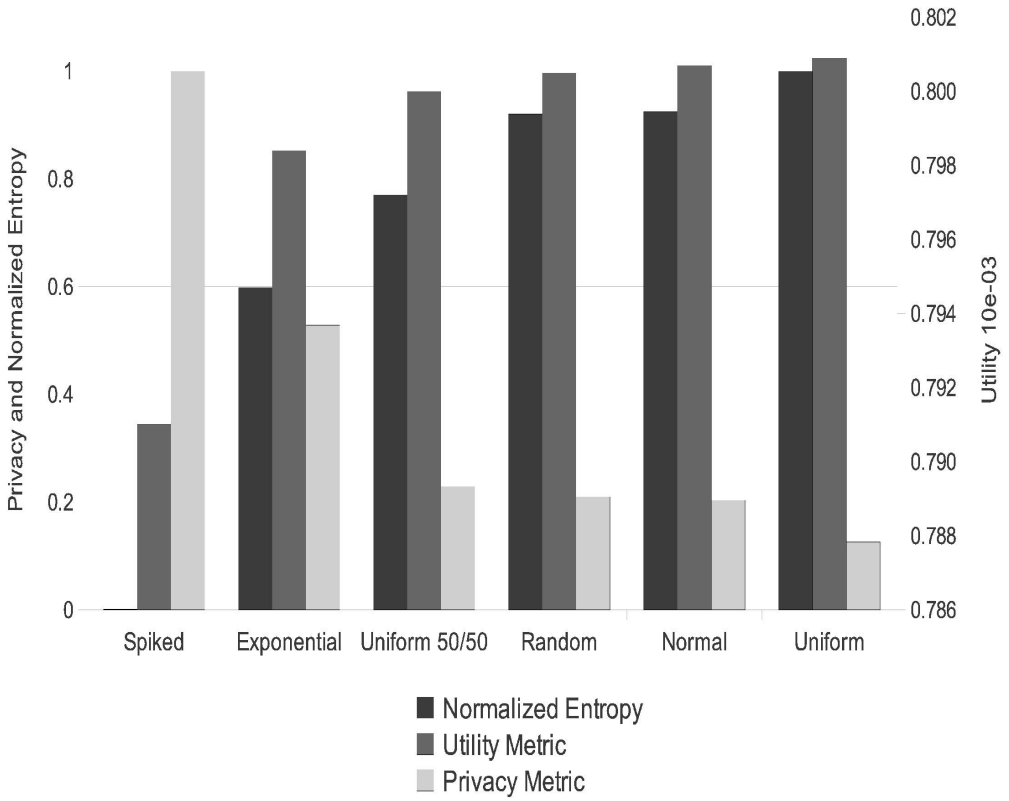

Variance and Covariance

- Variance

$$Var(\vec{x}^i) = \frac{\left(\left[\prod_{i=1}^{D} \alpha_i\right] - 1\right)^2}{N} \cdot \mathrm{P}\left(\mathrm{Y} = \vec{x}^i\right) \cdot \left(1 - \mathrm{P}\left(\mathrm{Y} = \vec{x}^i\right)\right)$$

- Covariance

$$Cov(\vec{x}^i, \vec{x}^j) = \frac{\left(\left[\prod_{i=1}^{D} \alpha_i\right] - 1\right)^2}{N} \cdot \mathrm{P}\left(\mathrm{Y} = \vec{x}^i\right) \cdot \mathrm{P}\left(\mathrm{Y} = \vec{x}^j\right)$$

- These equations produce the same results when compared to the Kronecker method.

Complexities of the Reconstruction Equations

- One-dimensional matrix multiplication:
  - $O(\alpha^2)$
- One-dimensional negative survey equation:
  - $O(\alpha)$
- *D*-dimensional reconstruction of natural extension.
  - $\mathrm{O}\left(\left[\prod_{i=1}^{D} \alpha_i\right] \cdot \left[\sum_{i=1}^{D} \binom{D}{i} (\mathrm{D} - \mathrm{i}) \cdot \alpha_{max}\right]\right)$
- *D*-dimensional reconstruction with matrixes.
  - $\mathrm{O}\left(\sum_{i=1}^{D} \left[\prod_{j=1, j \neq i}^{D} \alpha_i^2 \alpha_j\right]\right)$
- *D*-dimensional reconstruction with negative surveys
  - $\mathrm{O}\left(\mathrm{D} \cdot \prod_{i=1}^{D} \alpha_i\right)$

Kronecker Technique

- Turns a MDNS into a SDNS.
  - The perturbation matrix used is the Kronecker Product of the individual perturbation matrices for each dimension.
  - Y is marshalled into a vector.
  - The kronecker product is multiplied with the vector to obtain the estimated distribution A.
  - A is then demarshalled.

MDNS Strengths

- Energy efficient at both node and base station.
- All samples are guaranteed to be preserved.
- Allows dimensional adjustment which other perturbation matrices do not.
- Allows utility to be known *apriori.*
- Most optimal warner scheme.

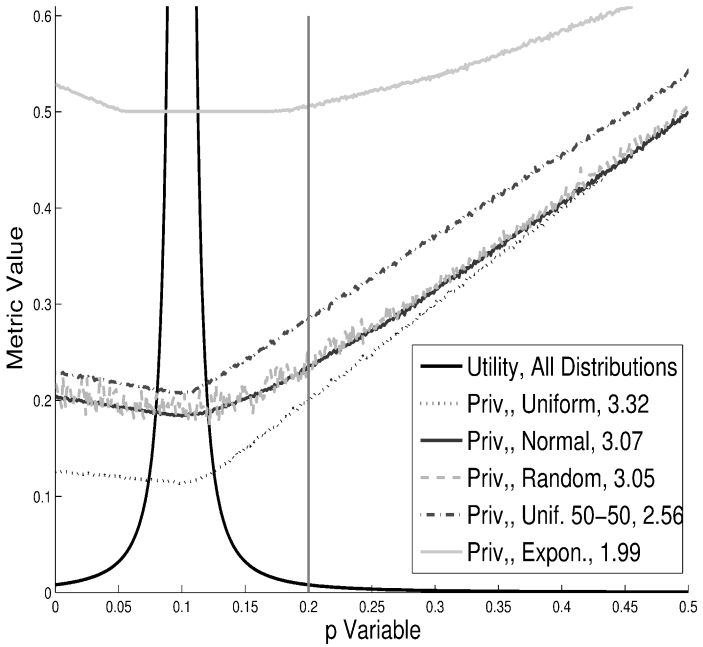

Magnification of Errors

- Probability of the expected value for the bins of Y deviating passed one standard deviation.