TriggerCast: Enabling Wireless Constructive Collisions

Yin Wang[†], Yuan He[†§], Dapeng Cheng[*], Yunhao Liu[†§], Xiang-yang Li[‡†]
[†]MOE Key Lab for Information System Security, School of Software, TNLIST, Tsinghua University
[§]Department of Computer Science and Engineering, HKUST
[*]Key Laboratory of Intelligent Information Processing in Universities of Shandong
(Shandong Institute of Business and Technology)
[‡]Department of Computer Science, Illinois Institute of Technology
{wangyin, he, yunhao}@greenorbs.com, chengdapeng@gmail.com, xli@cs.iit.edu

***Abstract*—Constructive Interference (CI) proposed in the existing work (e.g., A-MAC [1], Glossy [2]) may degrade the packet reception performance in terms of Packet Reception Ratio (PRR) and Received Signal Strength Indication (RSSI). The packet reception performance of a set of nodes transmitting simultaneously might be no better than that of any single node transmitting individually. In this paper, we redefine CI and propose TriggerCast, a practical wireless architecture which ensures concurrent transmissions of an identical packet to interfere constructively rather than to interfere non-destructively. CI potentially allows orders of magnitude reductions in energy consumption and improvements in link quality. Moreover, we *for the first time* present a theoretical sufficient condition for generating CI with IEEE 802.15.4 radio: concurrent transmissions with an identical packet should be synchronized at chip level. Meanwhile, co-senders participating in concurrent transmissions should be carefully selected, and the starting instants for the concurrent transmissions should be aligned. Based on the sufficient condition, we propose practical techniques to effectively compensate propagation and radio processing delays. TriggerCast has 95^{th} percentile synchronization errors of at most 250ns. Extensive experiments in practical testbeds reveal that TriggerCast significantly improves PRR (from 5% to 70% with 7 concurrent senders, from 50% to 98.3% with 6 senders) and RSSI (about 6dB with 5 senders).**

I. INTRODUCTION

In wireless Sensor Networks (WSNs), it is widely accepted that simultaneous transmissions will result in packet collisions. Recently, A-MAC [1] and Glossy [2] show that it is feasible for a common receiver to decode concurrent transmissions of an identical packet with high probability, if multiple transmissions are accurately synchronized. Their works basically operate on the passive side. In other words, they enable simultaneous transmissions to interfere non-destructively, namely to generate Non-Destructive Interference (NDI), in order to enhance network concurrency. Unfortunately, the packet reception performance of NDI might be no better than that of any single node transmitting individually (Fig. 1(a)), indicating NDI might degrade the performance of packet reception.

Our work advances the technique by actively utilizing the capacity of Constructive Interference (CI) to potentially improve the received power and link quality (Fig. 1(b)). CI is especially attractive for WSNs, because it potentially improves energy efficiency, and thus mitigates the limited power supply issue. A set of N nodes can achieve an N^2-fold increase in the received power of *baseband* signals, compared to a single node transmitting individually. It indicates that, to achieve the same Signal Noise Ratio (SNR), each node can reduce signal power with a factor of $\frac{1}{N^2}$, and the total power consumed by N nodes can be $\frac{1}{N}$ of the power required by a single sender. Moreover, simultaneously forwarding a packet can harness signal superposition gain, to improve Received Signal Strength Indication (RSSI) and Packet Reception Ratio (PRR).

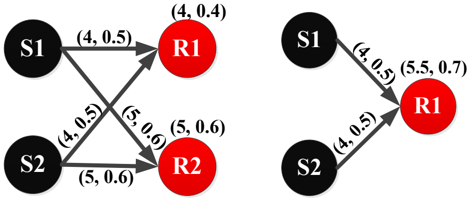

(a) Non-destructive interference (b) Constructive interference

Fig. 1. Both NDI and CI enable concurrency. Only CI improves RSSI and PRR. Here, we use (a, b) to describe a link, while a and b represent the RSSI and PRR respectively.

However, implementing CI in WSNs is challenging due to the following reasons. First, simultaneous transmissions must be synchronized at the chip level, namely 0.5μs for IEEE 802.15.4 radio. To generate NDI, Glossy's synchronization is sufficient, since it compensates most factors, such as clock drifts, software routine uncertainties of OS as well as asynchronous clocks (e.g., transmitter's radio and receiver's radio, MCU and radio module). However, it is not sufficient to construct CI. The Propagation delays and the radio processing delays significantly influence CI generation. Even worse, estimating the radio processing delays is an especially challenging task, as it varies from packet to packet, depends on the SNR, and is affected by the channel. Besides, in the absence of a central controller or a shared clock (e.g., GPS), they can only rely on their own radio signals as references.

Second, even if simultaneous transmissions are perfectly synchronized, i.e. no phase offset, they might not guarantee CI. The reason is because a radio signal has noise. Although signals are exactly aligned, noises also superpose. Whether SNR of the combined signal increases depends on SNRs and the Tx powers of individual signals.

Third, sensor nodes are always battery-powered, and have limited computational resources. It is difficult or even impossible to deploy complex signal processing algorithms in

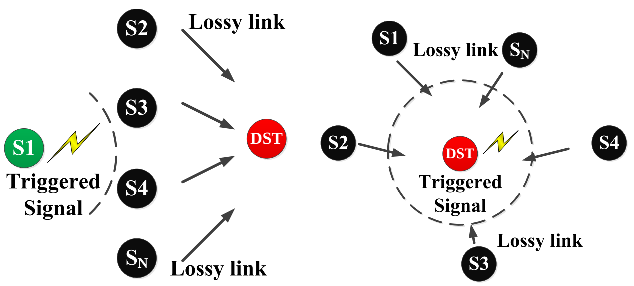

(a) sender-initiated (b) receiver-initiated

Fig. 2. TriggerCast: a radio triggered concurrent transmission architecture.

Commercial Off-The-Shelf (COTS) sensor platforms.

We propose TriggerCast, a practical distributed middleware to generate CI in WSNs. TriggerCast enables a co-sender (sender-initiated TriggerCast, Fig. 2(a)) or a receiver (receiver-initiated TriggerCast, Fig. 2(b)) to trigger a radio signal, which acts as a common reference for all concurrent senders to implement synchronized transmissions. TriggerCast can be used to control network topology without increasing Tx power of any node or adding new nodes. Several disconnected links can be controlled to form a connected link, which opportunistically reduces latency of routing (Fig. 3(a)). TriggerCast can also reduce packet retransmissions by improving PRR (Fig. 3(b)).

There are two key modules implemented in TriggerCast, namely, the Chip Level Synchronization (CLS) and the Link Selection and Alignment (LSA) algorithm. CLS enables concurrent transmissions to be synchronized within 0.5μs, by compensating the propagation and radio processing delays. Our experiments demonstrate that CLS has 95th percentile synchronization errors of at most 250ns. The accuracy is bounded by the running frequency (4,194,304Hz) of on-board MCU of TMote Sky sensor node. TriggerCast's LSA algorithm intelligently decides which co-senders to participate in simultaneous transmissions, and aligns their transmission time to maximize the overall link PRR under the condition of maximum system robustness. The underling CLS and LSA algorithms together ensure TriggerCast to generate CI in a practical testbed. Extensive experiments show that TriggerCast improves PRR (from 5% to 70% with 7 concurrent senders, and from 50% to 98.3% with 6 senders) and RSSI (about 6dB with 5 senders).

The contributions of this paper are summarized as follows.

i) We are the *first* to provide a *theoretical sufficient condition* for generating CI in WSNs.
ii) We propose TriggerCast, a practical middleware to ensure concurrent transmissions to interfere constructively rather than non-destructively. TriggerCast has 95th percentile synchronization errors of at most 250ns.
iii) The implementation of TriggerCast adopts COTS sensor nodes. According to the experiment results, the performance gains brought by TriggerCast are convincing.

II. RELATED WORK

Exploiting concurrent transmissions while suppressing interference is a promising direction, for its ability to decode packets from collisions, increase network throughput [3], improve power efficiency [4], enhance packet transmission reliability

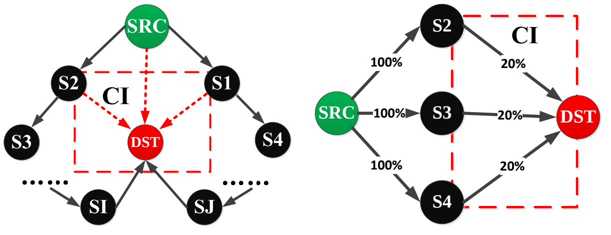

(a) Reduce transfer latency (b) Reduce data retransmissions

Fig. 3. (a) TriggerCast generates a new link from three disconnected links and thus reduce data forwarding latency. (b) TriggerCast makes use of signal superposition to improve PRR and hence reduce retransmission times.

[5] [6], and reduce latency of flooding [7] and data collection [8]. Prior works can be categorized as signal processing based and physical-layer phenomenon based.

Works based on signal processing include SIC [9] and Zigzag [10] for interference cancelation, 802.11n+ [11] for interference alignment in MIMO, AutoMAC [12] for rateless coding, and Simple Rule for chip error pattern [13]. Unfortunately, these signal processing algorithms cannot be directly applied in WSNs, in which sensor nodes have insufficient computation resources and limited energy supplies.

Physical-layer phenomenon based works mainly focus on exploring wireless radio properties of COTS transceivers. Such physical-layer phenomena mainly include capture effect [14] and message-in-message (MIM) [15]. Capture effect requires that the Signal Of Interest (SOI) is sufficiently stronger than the sum of the interfering signals. MIM needs special hardware support to continuously synchronize with the preamble of stronger signal. Both capture effect and MIM can only decode the stronger signal at the cost of dropping the other signals.

Recently, Backcast [1] experimentally discovers that, concurrent transmissions of short acknowledgment packets automatically generated by the radio hardware can interfere non-destructively. By implementing elaborate and accurate timing controls, Glossy [2] and SCIF [14] advances NDI of data packets instead of acknowledgment packets. However, the main purpose of our work is to make concurrent transmissions of an identical packet interfere constructively.

III. A SUFFICIENT CONDITION FOR GENERATING CI

The basic principle of 802.15.4 PHY layer is elaborated in [16]. We suppose there are N transmitters $\{T_i, i = 1, 2, ..., N\}$ simultaneously sending a same packet to a common receiver R. The output signal from each transmitter T_i arriving at the antenna of the receiver R is denoted as $S_R^i(t)$. The received signal $S_R(t)$ can be expressed as

$$S_R^i(t) = S_{msk}^i(t) * H^i(t) + N^i(t), \qquad (1)$$

where $S_{msk}^i(t)$ is the ith transmitted signal after MSK modulation, $H^i(t)$ and $N^i(t)$ denote the corresponding channel response and noise respectively. The received superposed signal $\overline{S_R(t)}$ is the sum of the N output signals $S_R^i(t)$. Hence we can approach

$$\overline{S_R(t)} = \sum_{i=1}^{N} (A_i S_R^i(t-\tau_i) + N_i(t)), |\tau_i| \le T_c \qquad (2)$$

where A_i and τ_i respectively depict the unified amplitude and phase offset of the ith arriving signal relative to the instant when the strongest signal reaches the receiver, $T_c(=0.5\mu)$s is the duration of a chip in IEEE 802.15.4 radio. Let λ_i be the SNR of the output signal $S_R^i(t)$, P_i denote average power of signal $S_R^i(t)$ and N_i represent power of noise $N_i(t)$. Obviously, we have $\lambda_i = \frac{P_i}{N_i}$. Let $S_R^1(t)$ be the strongest signal. Therefore, we have $A_1 = 1$, $\tau_1 = 0$, $P_i = P_1 {A_i}^2$. According to [14], it can be derived that the effective power $\overline{P}$ of superposed signals after demodulation is $\bar{P} = P_1(\sum_{i=1}^{N} A_i \cos(\omega_c \tau_i))^2$, while the aggregated power of noise $\overline{N}$ is $\sum_{i=1}^{N} \frac{P_i}{\lambda_i}$. As a result, the SNR of the received superposed signal is

$$\frac{\overline{P}}{\overline{N}} = \frac{P_1(\sum\limits_{i=1}^{N} A_i \cos(\omega_c \tau_i))^2}{\sum\limits_{i=1}^{N} P_i/\lambda_i} \leq \frac{P_1 \sum\limits_{i=1}^{N} A_i^2 \sum\limits_{i=1}^{N} (\cos(\omega_c \tau_i))^2}{P_1 \sum\limits_{i=1}^{N} A_i^2/\lambda_i}. \quad (3)$$

The inequality (3) can be derived by Cauchy-Schwarz inequality and equality holds if the condition satisfies

$$\frac{A_i}{\cos(\omega_c \tau_i)} = \frac{A_j}{\cos(\omega_c \tau_j)}, \qquad (\forall i, j). \quad (4)$$

To guarantee the received SNR of the superposed signal is better than the SNR of any single signal in the worst case, namely to ensure simultaneous transmissions to interfere positively, it is required that the maximum value of the received SNR is no less than $\lambda_{\max}$

$$(\frac{\overline{P}}{\overline{N}})_{max} > \lambda_{\min} \sum_{i=1}^{N} (\cos(\omega_c \tau_i))^2 \geq \lambda_{\max}. \quad (5)$$

Consequently, we derive a theoretical *sufficient condition (SC)* for CI with IEEE 802.15.4 radio.

i) Concurrent transmissions with a same packet should be synchronized at chip level, namely less than T_c=0.5μs;
ii) The phase offset of the ith arriving signal should satisfy: $|\tau_i| \leq \cos^{-1}(\sqrt{\frac{P_i}{P_1}}/\omega_c)$ (SC-I);
iii) The ratio of the minimum SNR $\lambda_{\min}$ and the maximum SNR $\lambda_{\max}$ of concurrent transmissions should satisfy: $\frac{\lambda_{\min}}{\lambda_{\max}} \geq \frac{1}{\sum_{i=1}^{N} (\cos(\omega_c \tau_i))^2}$ (SC-II).

IV. TRIGGERCAST IMPLEMENTATION

A. Chip Level Synchronization (CLS)

Eliminating the propagation delays and the radio processing delays in realistic environment is very challenging. Those delays vary from one packet to another, and are influenced by communication link qualities, asynchronous radio clocks, clock drifts as well as quantization errors. Fortunately, according to the *law of large numbers*, we can obtain the expected propagation and radio processing delays by a large number of trials. We select one transmitter-receiver pair which is 40 meters away in a indoor environment, and let the transmitter periodically send a packet every 500ms. Once the receiver successfully decodes a packet, it piggybacks a reply as soon as possible to the previous transmitter. As shown in Fig. 4, the time-stamps T_{S1} and T_{S2} represent the phases when the sender's radio starts transmitting a packet and ends a packet transmission, while the time-stamp T_{S3} denotes the phase when the radio begins a packet reception. The time-stamps T_{R1}, T_{R2} and T_{R3} characterize the phases when the receiver's radio starts a packet reception, ends a packet reception as well as begins a packet transmission respectively. The TMote Sky node can accurately capture the exact instants when MCU detects rising edge and falling edge of SFD interrupts, with MCU's timer capture functionality. The nth packet sent by the receiver includes time-stamps $T_{R1}(n)$, $T_{R2}(n)$ and $T_{R3}(n-1)$, which can be used by the transmitter, to evaluate the expected value of radio processing delay and propagation delay

$$\widehat{\Delta} = \frac{(\widehat{T_{S3}} - \widehat{T_{S1}}) - (\widehat{T_{R3}} - \widehat{T_{R1}})}{2}, \quad (6)$$

where the symbol $\widehat{\lambda}$ defines the mean value of λ.

Experimental results of delay measurement using Eq. (6) is displayed in Fig. 5 as the 'raw' curve. Unfortunately, the result is not sufficiently accurate. The measured delay ranges from 0.596μs to 5.01μs, with average value 2.32μs and variance 0.628μs. The instability of measured delay indicates that it is difficult to synchronize different transmitters at a magnitude of 0.5μs, if we straightly use the measured data for compensation. Fortunately, we disclose the data transmission delay is the same for all nodes. And thus we have $T_{S2}(n) - T_{S1}(n) = T_{R2}(n) - T_{R1}(n)$. The data transmission delays of the transmitter and the receiver are drawn in Fig. 6.

We also find that the measured data transmission delays are not stable for the transmitter-receiver pair. The reason for the instability is because of the jitters, clock drifts as well as hardware diversities of the nodes' DCOs. The drifts can be as high as 5000ppm in our measurement. We define $\chi(n) = (T_{S2}(n) - T_{S1}(n))/(T_{R2}(n) - T_{R1}(n))$ as the unified clock drift coefficient relative to the receiver. Consequently, we can calibrate Eq. (6) as

$$\widehat{\Delta}_{cal} = \frac{(\frac{T_{S3}\widehat{(n)-T_{S1}}(n)}{\chi(n)}) - (\widehat{T_{R3}} - \widehat{T_{R1}})}{2}. \quad (7)$$

We obtain the expected radio processing and propagation delay represented by DCO Ticks after the calibration of Eq. (7). To translate them to time, we also utilize the *Virtual High-resolution Time (VHT)* [17] approach, which calibrates the receiver's DCO with more stable external 32,768 Hz crystal as a reference. The measured propagation and radio precessing delay after clock drift calibration is shown as the 'drift calibration' curve in Fig. 5. The calibrated delay ranges from 3.66μs to 4.12μs, with average value 3.90μs and variance 0.012μs. We disclose that, in our measurements, the delays don't change so much as thought before. The measurement delay are almost constant, unless the nodes move or the channel significantly changes.

B. Link Selection and Alignment (LSA)

Assuming all the concurrent transmissions are synchronized at the chip level with CLS, according to the proposed sufficient condition in Section III, the problem to make concurrent transmissions superpose constructively can be formalized as CI-generation problem.

Problem: Let $\Phi = \{L_1, L_2, ..., L_N, L_i = (P_i, \lambda_i)\}$ define a lossy link set, where P_i and λ_i denote the received signal's RSSI and SNR of transmitter T_i respectively. The problem is to find a

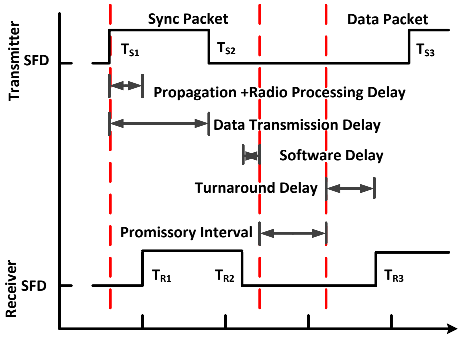

Fig. 4. TriggerCast's timing diagram of SFD signal for TMote Sky node (Without Preamble).

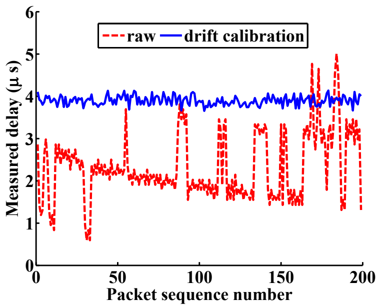

Fig. 5. Measured and calibrated delays of propagation and radio processing.

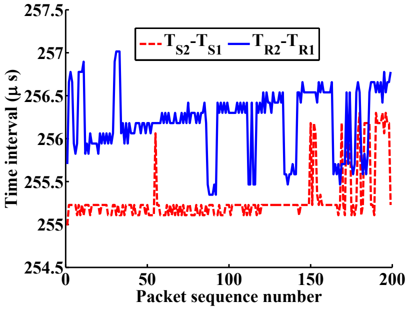

Fig. 6. Measured delays of data transmission for the transmitter and the receiver using DCO ticks.

Algorithm 1: Link Selection and Alignment

Input: Given a lossy link set $\Phi < P_i, \lambda_i >$. All link pairs of Φ are ordered.

Output: A lossy link subset $\Omega < P_j, \lambda_j, \tau_j >$ to maximize the superposed signal's SNR, where τ_j is the maximum allowed phase offset.

```
1  Sort $\Phi$ and store the result as $\Phi'$
2  Get the best link $< P, \lambda >$ in $\Phi'$ and insert $< P, \lambda, 0 >$ in
   empty set $\Omega$
3  for $i = 2 : N$ do
4      get link $< P_i, \lambda_i >$ in $\Phi'$;
5      calculate maximum allowed phase offset $\tau_i$ of link
       $< P_t, \lambda_t >$ using link $< P, \lambda >$, to satisfy (SC-I);
6      use set $\Omega$, link $< P_i, \lambda_i >$, phase offset $\tau_i$ to verify
       SC-II of the sufficient condition;
7      if step 6 satisfy (SC-II) then
8          insert link $< P_i, \lambda_i >$, phase offset $\tau_i$ to set $\Omega$;
9      else
10         break;
11     end
12     end
13 end
```

link subset Ω, which maximizes the superposed signal's SNR on condition that the combined link is better than any lossy link in Φ and the phase offset τ_i is as large as possible.

We define a link pair (L_i, L_j) is *ordered* if $P_i \geq P_j$ indicates $\lambda_i \geq \lambda_j$, $1 \leq i, j \leq N$. According to the sufficient condition for CI, it can be proved that this problem is NP-hard if there exists any disordered link pair in Φ. In practice, it is reasonable to assume all the link pairs in Φ are ordered. Generally speaking, for IEEE 802.15.4 signals, RSSI functions monotonically with PRR and thus with SNR. The pseudocode of LSA is described in algorithm 1. The time complexity of LSA algorithm is dominated by the sort function. Thus the time complexity of LSA algorithm is $O(nlog\ n)$.

Total compensation time: Consequently, we select the value of τ_i, to minimize the total number N_{com} of NOPs for the co-sender T_i

$$N_{com} = \left[(T - \widehat{\Delta}_{cal} + \tau_i) f_p \right]. \qquad (8)$$

where $[]$ is the round function, and T is a predefined maximum delay calibration time.

V. PERFORMANCE

We have implemented a prototype TriggerCast on TMote sky sensor nodes. The software is based on Contiki OS. During the overall TriggerCast's duration, except for the promissory interval, all the relevant interrupts and hardware timers that are not essential to TriggerCast's functioning are disabled. Since this interval is very short (several milliseconds), it is feasible that TriggerCast doesn't influence the upper layer's functionality. A runtime parameter adjustment software is developed, to make sure we can online change the system running parameters, without altering communication channels by programming the nodes.

A. Synchronization Accuracy

We first test the synchronization performance of multiple concurrent transmitters. We use three TMote sky nodes, one as a receiver and two as transmitters. We set the promissory interval parameter to 0 in receiver-initiated TriggerCast. We connect the SFD pins of the receiver (R) and one of the transmitters (S1) to a Agilent MSO-X serial oscilloscope. The other transmitter (S2) is 30 meters away in an indoor environment. However, it is difficult to measure the synchronization of S1 and S2 directly with the oscilloscope. As a result, we use R as a reference node. The synchronization between S1 and R can be monitored by the oscilloscope with a granularity of 5ns. The durations between T_{R1} and T_{R3} (Fig. 4) of the receiver are accurately measured when S1 and S2 transmit independently. The differences of the durations can be used for synchronization accuracy measurement, since both S1 and S2 rely on the instant T_{R1} as a reference. The CDF of synchronization errors compared with the Glossy synchronization algorithm is illustrated in Fig. 7. TriggerCast's CLS algorithm can synchronize multiple transmitters at a magnitude of 250ns. The accuracy is limited by the operating frequency of the MCU of TMote Sky sensor nodes. The Glossy synchronization algorithm degrades as the distance differences between two transmitter-receiver pairs increase. CLS outperforms Glossy because CLS compensates the time due to propagation and radio processing delays.

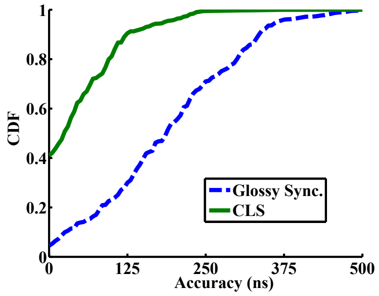

Fig. 7. Synchronization error of TriggerCast less than 250 ns has more than 95% confidence.

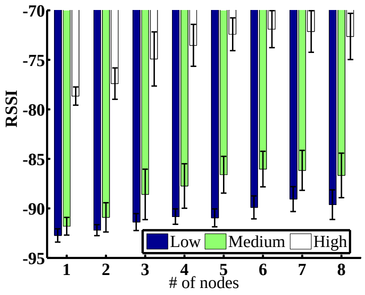

Fig. 8. TriggerCast increases RSSI.

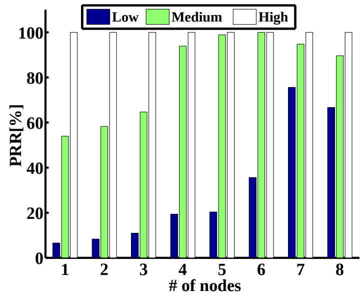

Fig. 9. TriggerCast improves PRR.

B. Power Gains and PRR Improvements

In practical experiments, up to 8 senders with all three different kinds of links (low (PRR $< 5\%$), medium ($5\% <$PRR$<$ 90%), high (PRR $> 90\%$)), are executed to transmit packets at the same time. Due to the limit of physical space, we randomly insert NOPs to simulate different propagation and radio processing delays. We adjust the received RSSIs of each sender's individual packet transmission to almost the same, to eliminate the influence of capture effect. All the results are averages of more than 1000 tests. Fig. 8 lists the power gains due to multiple senders of different link types (disconnected link: 1-5 dB, intermediate link: 2-6 dB, connected link: 2-6 dB). The maximum power gain can approach N^2 for N concurrent transmitters. Fig. 9 shows that PRR can be significantly improved by leveraging CI. For 7 disconnected links, the PRR achieves almost 70%, which is better than our previous understandings of harnessing sender diversity gain $(1-(1-0.05)^3 \approx 30.2\%)$. TriggerCast improves the PRR of intermediate links from 50% to almost 100% with 6 concurrent senders. Our experiments indicate that TriggerCast can *control network topology* (e.g., increasing new communication links) *without changing the original network state* (adding new nodes, increasing nodes' power, etc.). This characteristic is attractive to improve routing performance (explained in Fig. 3). To the best of our knowledge, we are the *first* to report that multiple concurrent transmitters can reach such PRR improvements in realistic WSNs.

VI. CONCLUSIONS AND FUTURE WORK

We propose TriggerCast, by far the first work to implement CI instead of NDI in WSNs, to the best of our knowledge. TriggerCast compensates the propagation and radio processing delays, conducts link selection, and ensures transmission alignment, so as to enable CI. We implement TriggerCast in real testbeds, and experimentally demonstrate the performance gains brought by TriggerCast. We also provide a theoretical sufficient condition on how to ensure concurrent transmissions interfere constructively. Our future work includes taking node mobility and low duty-cycle factors into account in TriggerCast, and exploiting applications of TriggerCast for time synchronization and localization.

ACKNOWLEDGMENT

This work is supported in part by NSFC program under Grant 61170213, 61170216, 61228202, 61070175, NSF CNS-0832120, NSF CNS-1035894, NSF ECCS-1247944, and China 973 Program under Grant No.2012CB316200 and No.2011CB302705, China 863 Program under grant No. 2011AA010100, China Postdoctoral Science Foundation under 2012T50096 and 2011M500019, NSF of Shandong under Grant No. ZR2011FL008, ZR2011FL004, Science and technology on communication information security control laboratory.

REFERENCES

[1] P. Dutta, S. Dawson-Haggerty, Y. Chen, C. Liang, and A. Terzis, "Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless," in *Proc. of ACM SenSys*, 2010.
[2] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with Glossy," in *Proc. of ACM/IEEE IPSN*, 2011.
[3] X. Wang, L. Fu, and C. Hu, "Multicast performance with hierarchical cooperation," *IEEE/ACM Trans. on Networking*, vol. 20, no. 3, pp. 917–930, 2010.
[4] Z. Cao, Y. He, and Y. Liu, "L^2: Lazy forwarding in low duty cycle wireless sensor networks," in *Proc. of IEEE INFOCOM*, 2012.
[5] Liu, W. and Liu, Y. and Zhong, Z. and Sun, L. and Zhu, H. and He, T., "Exploiting Ephemeral Link Correlation for Mobile Wireless Networks," in *Proc. of ACM SenSys*, 2012.
[6] J. Wang, Y. Liu, M. Li, W. Dong, and Y. He, "QoF: Towards comprehensive path quality measurement in wireless sensor networks," in *Proc. of IEEE INFOCOM*, 2011.
[7] S. Guo, Y. Gu, B. Jiang, and T. He, "Opportunistic flooding in low-duty-cycle wireless sensor networks with unreliable links," in *Proc. of ACM MobiCom*, 2010.
[8] Z. Li, M. Li, J. Wang, and Z. Cao, "Ubiquitous data collection for mobile users in wireless sensor networks," in *Proc. of IEEE INFOCOM*, 2011.
[9] D. Halperin, T. Anderson, and D. Wetherall, "Taking the sting out of carrier sense: interference cancellation for wireless lans," in *Proc. of ACM MOBICOM*, 2008.
[10] S. Gollakota and D. Katabi, "Zigzag decoding: combating hidden terminals in wireless networks," in *Proc. of ACM SIGCOMM*, 2008.
[11] K. Lin, S. Gollakota, and D. Katabi, "Random access heterogeneous mimo networks," in *Proc. of ACM SIGCOMM*, 2011.
[12] A. Gudipati, S. Perreira, and S. Katti, "AutoMAC: Rateless wireless concurrent medium access," in *Proc. of ACM MOBICOM*, 2012.
[13] K. Wu, H. Tan, H. Ngan, Y. Liu, and L. Ni, "Chip error pattern analysis in IEEE 802.15. 4," *IEEE Trans. on Mobile Computing*, vol. 11, no. 4, pp. 543–552, 2012.
[14] Y. Wang, Y. He, X. Mao, Y. Liu, and X. Li, "Exploiting constructive interference for scalable flooding in wireless networks," *to appear in IEEE/ACM Transactions on Networking (TON)*, 2012.
[15] N. Santhapuri, J. Manweiler, S. Sen, R. Choudhury, S. Nelakuditi, and K. Munagala, "Message in message (MIM): A case for reordering transmissions in wireless networks," in *Proc. of ACM HotNets-VII*, 2008.
[16] N. Oh and S. Lee, "Building a 2.4-GHz radio transceiver using ieee 802.15.4," *IEEE Circuits and Devices Magazine*, vol. 21, no. 6, pp. 43–51, 2006.
[17] T. Schmid, P. Dutta, and M. B. Srivastava, "High-resolution, low-power time synchronization an oxymoron no more," in *Proc. of ACM/IEEE IPSN*, 2010.