Critical Delay-Tolerant NB-IoT Applications in GNSS-Denied Direct-to-Satellite Environments

Miguel José Filipe
LAAS-CNRS, Université de Toulouse, CNRS
Toulouse, France
miguel.filipe@laas.fr

Nicola Accettura
LAAS-CNRS, Université de Toulouse, CNRS
Toulouse, France
nicola.accettura@laas.fr

***Abstract*—The increasing availability of cellular-enabled Low Earth Orbit (LEO) satellites is an opportunity to broaden the IoT market, and has already catalyzed standardization efforts by 3GPP regarding the Narrowband Internet of Things (NB-IoT) technology. Nonetheless, its correct functioning over Direct-to-Satellite links is impaired by the time-variation of both Doppler shift and propagation delay. To cope with that, 3GPP mandates the use of Global Navigation Satellite Systems (GNSS). Yet, this technological dependency is ill-suited to GNSS-denied scenarios. Previous attempts to enable GNSS-free Direct-to-Satellite NB-IoT (DtS NB-IoT) communications were all optimized for loss-tolerant applications requiring very small traffic rates. For the first time, this paper targets the collection of critical delay-tolerant sensor data through DtS NB-IoT, achievable through stable long-duration uplink connectivity: data are buffered during inter-pass intervals and dispatched over the whole duration of the next satellite pass. To do so, ground low power devices measure the frequency offset of downlink configuration signals to estimate the Doppler shift evolution, without dedicated hardware. At the cost of some additional computation, the available communication time in GNSS-denied environments is sufficiently extended to enable critical data delivery.**

***Index Terms*—IoT; NB-IoT; LEO; 3GPP; Direct-to-Satellite**

I. Introduction

Cellular networks cover roughly 90% of the world's population but only about 15% of the Earth's surface [1], leaving oceans, deserts and polar regions without terrestrial Internet of Things (IoT) connectivity. Low Earth Orbit (LEO) satellites remove that constraint, and Low Power Wide Area Networks have been identified as the key enablers of Direct-to-Satellite IoT [2]. Among them, Narrowband IoT (NB-IoT) is the most promising candidate: (i) its synchronized access implies a listen-first policy, which is more energy-efficient than a transmit-first strategy; (ii) its licensed spectrum limits interference that would otherwise erode the signal-to-noise ratio (SNR). The considered architecture includes NB-IoT User Equipments (UE) on the ground transmitting sensor data towards a LEO satellite equipped with an NB-IoT evolved NodeB (eNB). However, the communication in such a configured scenario suffers from time-varying Doppler shift and propagation delay. The relative motion of the satellite, at roughly 7.5 km/s, shifts the carrier frequency by tens of kilohertz and lengthens the propagation delay, so the ground equipment must predict and pre-compensate both quantities. 3GPP Release 17 [3] introduced normative Non-Terrestrial Network (NTN) support on GNSS-derived pre-compensation, which introduces a dependency that fails under jamming, and clashes with the energy budget of low-power devices. Several attempts tried to remove this dependency, while differing in which network element is chosen to perform compensation. In fact, when compensation is delegated to the satellite, the majority of the pre-compensation is computed by the eNB, leaving only a residual CFO for the UE to correct. This can be done through either quasi-Earth-fixed cells and narrow beams [4], or Earth-moving footprint with the timing advance being broadcast [5]. Both solutions require the deployment of more sophisticated and expensive satellite eNB. Instead, in other works the compensation burden is shared between eNB and UE. In [6] a two-carrier acquisition is used to separate the UE clock drift from the Doppler, while in [7] an instantaneous pre-compensation of the residual CFO is derived before any transmission; in both cases a more expensive satellite is required in addition to additional computational capabilities on the UE. Eventually, some works place the compensation burden only on the UE. According to [8], a UE can self-localize by time-difference-of-arrival over multiple Narrowband Primary Synchronization signal (NPSS), and then rely on the broadcast ephemeris for pre-compensation. Remarkably, none of the above information is required by [9]: the frequency offset measured across a cluster of NPSS yields the maximum elevation angle through least-squares approximation, from which the whole Doppler and delay curves can be inferred. Reconstructing the trajectory rather than tracking the instantaneous offset is what makes forward pre-compensation possible, thus requiring only a firmware update. That solution is nonetheless dimensioned for short non-critical transmissions: with a synchronization and random-access overhead of about 5.56 s [10], a UE communicating for 12 s [9] spends nearly half of its time budget preparing to transmit, against 15% as discussed in the solution presented hereafter. Building upon [9], the goal of this contribution is to enable the transmission of critical delay-tolerant sensor data: cargo tracking across oceans, livestock monitoring, leak detection along isolated pipelines, and glacier-melt stations on Mount Everest, all beyond the reach of terrestrial infrastructure. The proposed method leverages the existing synchronization process to measure the Carrier Frequency Offset (CFO) through bit-exact regeneration of received downlink configuration signals. Each regenerated block serves as a local reference, and

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

the resulting observations refine the least-squares estimation of the satellite trajectory. Unlike other GNSS-free approaches, which rely on standard modifications or hardware changes, or operate only at small uplink data rates, the method requires no change to the air interface and lets the UE exploit a larger part of the visibility window. The maximum communication window is extended to at least 37 s whenever the first NPSS is received within the first 66% of the pass. The remainder of this paper is organized as follows. After introducing the technological background in Sec. II, the least-squares formulation of [9] is further refined, as detailed in Sec. III. Then, the whole strategy is evaluated through extensive Monte-Carlo simulations in Sec. IV. Finally, Sec. V draws conclusions and envisages future works.

II. Technological Background

This section details the organization of the NB-IoT physical layer and downlink synchronization process [11]. Specifically, time is divided into 10 ms radio frames (RF), each partitioned into 10 subframes of 1 ms (SF#0 to SF#9). The decoding process starts by listening for an NPSS, transmitted in SF#5 of every RF. The NPSS provides symbol timing and a coarse CFO estimate; at least 1 is required to proceed. A cluster of NPSS is required to form the Doppler curve estimation of [9] before decoding the remaining signals. The next step is to decode the Narrowband Secondary Synchronization Signal (NSSS), which is placed in SF#9 of every even RF (i.e., every 20 ms) and resolves the cell identity, from which the Master Information Block (MIB-NB) is decoded. The MIB-NB is carried over a 640 ms period, split into 8 independently decodable 80 ms sequences, each repeating the same 1 ms payload 8 times in SF#0 of every RF. The MIB-NB also specifies the repetition count and resources allocated to the System Information Block Type 1 (SIB1-NB). The SIB1-NB occupies SF#4 of every even RF and recurs every 2560 ms, repeated 4, 8 or 16 times according to the coverage level; the maximum number of repetitions is adopted here. The SIB1-NB in turn schedules the System Information Block Type 2 (SIB2-NB) through 4 parameters that the 3GPP standard leaves to the deployment. The configuration adopted here delivers 80 repetitions of 8 downlink subframes each, within a 1600 ms System Information window (SI-window) recurring every 5120 ms.

The acquisition duration from NPSS to SIB2-NB is governed by the operating SNR, which varies with the UE position within the satellite beam [12] and the channel impairments. A lower SNR likely requires more repetitions per block and degrades the CFO estimate. Two regimes are therefore considered (High-SNR and Low-SNR), each characterized by the CFO estimation error standard deviation of the signal. Detecting 11 NPSS is assumed to take 1 s in both regimes [9]. At low SNR the initial synchronization may occupy up to 10 RF [13], so that after an NPSS, 5 repetitions may be needed for the NSSS, whereas a single one suffices at high SNR. For subsequent signals the duration also depends on frame alignment: a block decoded just after the next block's scheduled occurrence forces the UE to wait. Each regime is therefore bounded by a fastest and a slowest alignment; Table I reports the values used throughout this paper, together with the resulting average timings.

TABLE I: NB-IoT signal acquisition parameters and mean acquisition duration per regime.

Signal	Period (s)	Repetitions		Mean duration (ms)	
		$\sigma_{\text{signal}}^{\text{High-SNR}}$	$\sigma_{\text{signal}}^{\text{Low-SNR}}$	$\sigma_{\text{signal}}^{\text{High-SNR}}$	$\sigma_{\text{signal}}^{\text{Low-SNR}}$
NSSS	0.02	1	5	9	89
MIB-NB	0.08	1	8	526	666
SIB1-NB	2.56	1	16	184	2594
SIB2-NB	5.12	1	80	1775	3355
Total from last NPSS				**3494**	**7704**

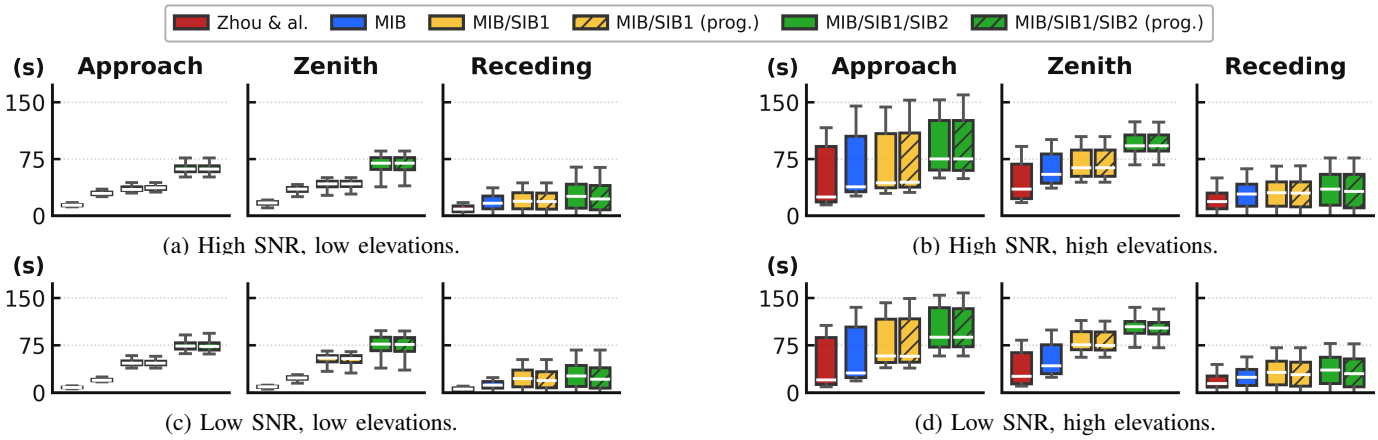
With the acquisition instants fixed, the remaining parameter is the CFO estimate precision per signal. The NSSS is not used in the CFO estimation, as its proximity to the NPSS cluster adds negligible information. For the remaining signals, the per-signal standard deviation is scaled from the NPSS baseline by the number of time-domain samples,

$$\sigma_{\text{signal}} = \sqrt{\frac{N_{\text{NPSS}}}{N_{\text{signal}}}} \cdot \sigma_{\text{NPSS}} \tag{1}$$

where σ_{NPSS} and N_{NPSS} are the CFO error standard deviation and the sample count for the NPSS, and σ_{signal}, N_{signal} the same for the signal considered. The scaling is conservative with respect to the Cramér–Rao lower bound (CRLB), whose variance decreases faster than linearly in the number of samples [14]: applying (1) therefore overestimates σ_{signal} for every signal beyond the NPSS, so the gains reported below are conservative and a tighter error model can only increase them. A rigorous treatment would bind each signal to its own CRLB and the SNR threshold above which it is attained. That analysis is left to future work; the present aim is to establish that the communication window can be extended. The two regimes are anchored at $\sigma_{\text{NPSS}}^{\text{High-SNR}} = 100$ Hz and $\sigma_{\text{NPSS}}^{\text{Low-SNR}} = 200$ Hz, respectively. With $N_{\text{signal}} = 1508$, 12068 and 15360 samples for the NPSS, MIB-NB and SIB1-NB/SIB2-NB, (1) gives $\sigma_{\text{signal}} = 35.3$ Hz and 31.3 Hz in the high-SNR regime, and 70.7 Hz and 62.7 Hz in the low-SNR one.

III. Proposed Approach

The approach builds on the least-squares Doppler estimation of [9]. Once the initial pre-compensation is applied, the UE decodes the following downlink signals as in Sec. II; each is then re-encoded into a local reference, and the CFO measured by correlating it against the received samples. Carrying more time-domain samples than the NPSS, these signals yield more precise observations at no additional reception cost, according to (1). In [9] the 11 NPSS fall within a short interval and share the same standard deviation, so a plain least-squares fit suffices. Instead, the additional signals are acquired up to 3.5 s later at high SNR and 7.7 s at low SNR (Table I), with a smaller CFO error standard deviation, so each estimate must be weighted by its relative precision. The Doppler curve $g(\theta^*, t)$ is a function of



(a) High SNR, low elevations.

(b) High SNR, high elevations.

(c) Low SNR, low elevations.

(d) Low SNR, high elevations.

Fig. 1: Distribution of maximum communication time by pass phase, for combinations of SNR level and elevation range.

time t and the maximum elevation angle θ^*. The least-squares adjusts the guessed $\hat{\theta}^*$ and timestamp $\hat{t}_0$ of the first NPSS to the measured values f_i at times $t_0 + \Delta t_i$, with i indicating the i-th CFO estimation after the very first NPSS. Under independent Gaussian noise, the estimator weights each residual by its inverse variance, so that a more precise observation contributes proportionally more to the estimate [15]:

$$(\hat{\theta}^*, \hat{t}_0) = \arg\min_{\theta^*,\, t_0} \sum_i \frac{\left[g(\theta^*, t_0 + \Delta t_i) - f_i\right]^2}{\sigma^2_{\text{signal}}}, \qquad (2)$$

where σ_{signal} denotes the standard deviation of the signal that produced the i-th f_i value at time $t_0 + \Delta t_i$.

Whenever several signals are decoded, all are used; the remaining choice is whether the estimate is computed in one shot or progressively. The batch approach waits until the eligible signals have been decoded and uses the full set at once. The progressive approach instead refines the estimate as each signal is decoded. Since the time available to decode a signal depends on the current estimate, often only one or two are reachable; refining early extends the attainable depth, at the risk that an intermediate update degrades the baseline.

IV. Simulation Results

The methods of Sec. III are evaluated through Monte-Carlo simulations. 11 NPSS are generated with their exact Doppler shifts, each corrupted by zero-mean Gaussian noise with the standard deviation of (1). A plain least-squares fit yields an initial estimate of θ^* and t_0, from which a first maximum communication duration follows, defined as the time over which the estimated curve stays within 500 Hz of the true one [9]. That duration determines which signals can be decoded in time; each contributes one further CFO measurement at its own σ_{signal}, and a weighted re-estimation (Sec. III) produces a new duration, compared against the initial prediction to quantify the gain. The simulation considers a candidate t_0 approximately every 2 s, and it runs the least-squares estimation 3000 times per t_0. This process is repeated for θ^* from 35° to 90° in 5° increments, and for both SNR regimes. The satellite pass is split in 3 equal phases, in respect to the given visibility window for a given θ^*: Zenith, Approach and Receding.

Fig. 1 reports the distribution of the maximum communication duration by SNR regime, elevation range (low: 35°-55°, high: 60°-90°) and pass phase. Adding signals improves the Doppler curve estimation in every combination. The MIB-NB alone already doubles the maximum communication time of [9]: 14 → 29 s in Approach, 17 → 35 s at Zenith and 9 → 17 s in Receding (Fig. 1a). High elevations yield longer windows, since low-elevation trajectories produce Doppler curves of almost constant negative slope and are therefore poorly identified, whereas high-elevation ones vary more sharply around the Zenith. The low-SNR regime outperforms the high-SNR one. While this result appears counter-intuitive, it follows from Sec. III: the information carried by an added observation grows with its distance from the NPSS cluster, and the longer acquisition of the SIB1-NB at low-SNR (cf. Table I) compensates its larger standard deviation. Decoding up to SIB2-NB preserves this ranking and further increases the maximum communication time. At low SNR and high elevation it grows from 58 to 88 s in Approach and from 76 to 105 s at Zenith, that is 4.4× and 4× the baselines of 20 s and 26 s (Fig. 1d). At low elevation the gain reaches 9×, with 8 → 74 s in Approach and 9 → 77 s at Zenith (Fig. 1c), while at high SNR the least favorable case still gains 2×, from 19 to 35 s in Receding (Fig. 1b). In the Receding phase the third block adds little, 32 → 36 s in Fig. 1d, so the additional computation is not justified. Fig. 2 reports the fraction of the remaining visibility time usable after the latest least-squares estimation: the baseline [9] exploits between 0.13 and 0.37 depending on regime and elevation, whereas decoding through the SIB2-NB raises it to between 0.64 and 0.72.

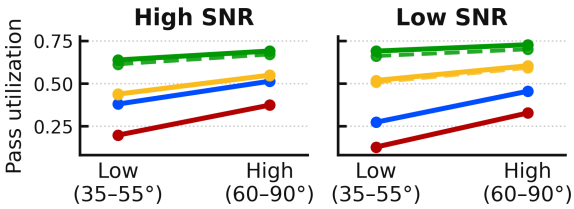
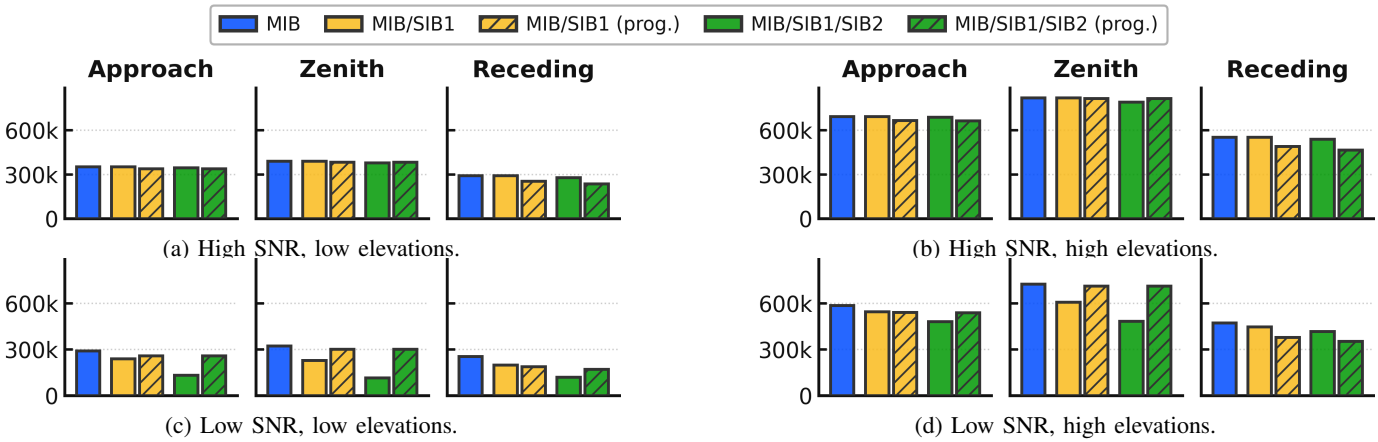


Fig. 2: Fraction of the pass remaining, averaged over pass phases, per SNR regime and elevation range.



(a) High SNR, low elevations.

(b) High SNR, high elevations.

(c) Low SNR, low elevations.

(d) Low SNR, high elevations.

Fig. 3: Distribution of the number of usable satellite passes by pass phase, for combinations of SNR level and elevation range.

Interestingly, the progressive and batch variants yield similar durations; they differ in how often the deeper signals are reachable at all (Fig. 3). An estimate built on the NPSS cluster alone may leave enough time for the MIB-NB but not for the SIB1-NB, whereas re-computing it once the MIB-NB is decoded makes the SIB1-NB reachable; the same mechanism then extends to the SIB2-NB. At low-SNR and low elevation the baseline offers only 6–9 s of communication time, which is frequently insufficient to reach the SIB2-NB. In addition, refining the estimate after each decoded signal roughly doubles the number of passes that reach SIB2-NB, i.e., $\approx$259k against $\approx$ 133k in the Approach phase (Fig. 3c). Therefore, the advantage is larger where the full set is harder to decode, and it costs at most 0.03 of the exploited pass fraction (Fig. 2). At high SNR the baseline already provides enough time to reach the last block, so the two variants coincide.

V. Conclusion

This paper presented a novel dynamic correction layer built on top of a lightweight frequency offset estimation strategy for Direct-to-Satellite NB-IoT communications with LEO satellites. Each decoded signal is regenerated bit-exactly and correlated against the received samples to yield an additional CFO observation, which is folded into a weighted least-squares re-estimation of the satellite trajectory. The maximum communication time reaches up to 105 s, that is up to 72% of the remaining pass, sufficient to dispatch buffered delay-tolerant data within a single pass. Future work will focus on experimental validation, on an in-depth study of the CRLB, on benchmarking the energy cost of bit-exact regeneration against a GNSS.

Acknowledgment

This research was funded by the French government under the France 2030 ANR program “PEPR Networks of the Future” through the DONUTS project (ref. ANR-24-PEFT-0001), and by the French National Research Agency (ANR) under the STEREO project (ref. ANR-22-CE25-0014).

References

[1] Eseye, “How IoT satellite communications revolutionize maritime operations,” https://www.eseye.com/resources/blogs/how-iot-satellite-communications-revolutionize-maritime-operations/, 2025, [Online].

[2] J. A. Fraire, S. Céspedes, and N. Accettura, “Direct-to-satellite IoT – a survey of the state of the art and future research perspectives,” in *Ad-Hoc, Mobile, and Wireless Networks (ADHOC-NOW 2019)*. Luxembourg, Luxembourg: Springer, Oct. 2019, pp. 241–258.

[3] 3GPP, “Study on NB-IoT / eMTC support for non-terrestrial networks (NTN),” 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 36.763, 2021, release 17, V17.0.0.

[4] M.-G. Kim and H.-S. Jo, “Performance analysis of NB-IoT uplink in low earth orbit non-terrestrial networks,” *Sensors*, vol. 22, no. 18, p. 7097, Sep. 2022.

[5] J.-C. Lin, “A unified random access procedure for an NB-IoT protocol in nonterrestrial and terrestrial networks,” *IEEE Internet of Things Journal*, vol. 12, no. 3, pp. 3424–3440, Feb. 2025.

[6] M. Conti, S. Andrenacci, N. Maturo, S. Chatzinotas, and A. Vanelli-Coralli, “Doppler impact analysis for NB-IoT and satellite systems integration,” in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, Jun. 2020, pp. 1–7.

[7] H. Chougrani, S. Kisseleff, W. A. Martins, and S. Chatzinotas, “NB-IoT random access for nonterrestrial networks: Preamble detection and uplink synchronization,” *IEEE Internet of Things Journal*, vol. 9, no. 16, pp. 14 913–14 927, Aug. 2022.

[8] Q. Qi, T. Hong, and G. Zhang, “Positioning-based uplink synchronization method for NB-IoT in LEO satellite networks,” *Symmetry*, vol. 17, no. 7, p. 984, Jul. 2025.

[9] Z. Zhou, N. Accettura, R. Prévost, and P. Berthou, “Lightweight synchronization to NB-IoT enabled LEO satellites through Doppler prediction,” in *Proc. 19th Int. Conf. Wireless Mobile Comput. Netw. Commun. (WiMob)*, Montreal, QC, Canada, Jun. 2023, pp. 218–223.

[10] N. Accurso, N. Mastronarde, and F. Malandra, “A comprehensive analytical model of end-to-end uplink delay in 5G NB-IoT networks,” in *Proc. 5th IEEE Int. Mediterranean Conf. Commun. Netw. (MeditCom)*, Nice, France, Jul. 2025, pp. 1–6.

[11] M. Kanj, V. Savaux, and M. Le Guen, “A tutorial on NB-IoT physical layer design,” *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2408–2446, 2020.

[12] S. Cluzel, L. Franck, J. Radzik, S. Cazalens, M. Dervin, C. Baudoin, and D. Dragomirescu, “3GPP NB-IoT coverage extension using LEO satellites,” in *Proc. IEEE 87th Veh. Technol. Conf. (VTC Spring)*, Porto, Portugal, Jun. 2018, pp. 1–5.

[13] R. B. Sørensen, H. K. Møller, and P. Koch, “5g NB-IoT via low density LEO constellations,” in *Proceedings of the AIAA/USU Conference on Small Satellites*, Logan, UT, USA, Aug. 2021, paper SSC21-X-03, 35th Annual Small Satellite Conference.

[14] D. C. Rife and R. R. Boorstyn, “Single tone parameter estimation from discrete-time observations,” *IEEE Trans. Inf. Theory*, vol. 20, no. 5, pp. 591–598, Sep. 1974.

[15] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., ser. Springer Series in Operations Research. New York: Springer, 2006.