CONTRIBUTED PAPER

Rescuing the Huygens Mission From Fiasco

When the space probe mission to Saturn's moon Titan was threatened, a joint European-American effort analyzed the radio receiver flaw and redesigned the relay-link geometry.

By Luitjens Popken, *Senior Member IEEE*

ABSTRACT | On January 14, 2005, the Huygens probe, which forms part of the joint NASA/ESA/ASI deep-space mission Cassini-Huygens, accomplished its spectacular descent through the atmosphere of Titan, Saturn's largest moon. This mission success, however, became possible only after the rescue of the endeavor from an implementation flaw that was discovered in 2000 during an in-orbit test of the Huygens relay-link receiver. The problem threatened to cripple the entire $350 million probe mission.

This paper presents the model of the faulty implementation of a data transition tracking loop. This model was the driver for the trajectory redesign leading to the revised relay-link geometry of the Huygens mission. The recovery eventually allowed for the data retrieval from the probe during its descent and after the landing on Titan, despite the tracking deficiency in the symbol synchronizer. The approach taken to identify and solve the communication problem is also an illustrative tutorial example of synchronization theory applied to save an entire space mission from disaster.

KEYWORDS | Cassini orbiter; data transition tracking loop (DTTL); Huygens probe; mission recovery; Saturn; synchronization; Titan

I. BACKGROUND AND INTRODUCTION

The Cassini spacecraft reached planet Saturn in July 2004. It had carried the Huygens space probe, which was released from the orbiter on Christmas Day 2004, sending it coasting towards a parachute landing on Titan, Saturn's largest moon. On January 14, 2005, the probe performed its fascinating three-hour descent through the dense atmosphere of Titan.[1]

About five years earlier, however, in February 2000, a significant noncompliance of the relay-link receiver had been detected during an in-orbit performance test [1], [2]. Additional performance mappings could only ascertain the receiver deficiency, which threatened to cripple the entire Huygens probe mission. It would have inflicted very substantial data losses if the foreseen original mission plan with the corresponding relay-link geometry would have been followed as intended. Unfortunately, the earlier stage of the ground testing and verifications had failed to detect the problem, due to test coverage limitations.

To attempt for a redesign of the probe mission [3], [4], it was crucial to identify what caused the technical anomaly to provide an accurate and verified model for the given receiver characteristic and to predict reliably the remaining performance. The flaw was tracked down to an implementation problem in the symbol synchronizer of the receiver. The circuit is incompatible with the time-varying relative relay geometry and the implied Doppler effect on the (fixed) data rate. This relay data rate is significantly higher than the maximum uplink rates commonly applied for spacecraft telecommanding.

The symbol synchronizer design is based upon the well-established concept of the data transition tracking loop (DTTL) [5]. However, in the actual implementation, the loop tracking performance is significantly degraded by a loop parameter setting that is incompatible with the original mission. In addition, the performance is further constrained by two mutually dependent automatic gain control (AGC) loops (noncoherent and coherent). The coherent AGC interacts directly with the DTTL so that the bandwidth and tracking capability of the synchronizer are discontinuous functions of the received signal power.

Manuscript received June 21, 2007. This work was presented in part at the 2004 IEEE Aerospace Conference Big Sky, MT, March 6–13, 2004.

Digital Object Identifier: 10.1109/JPROC.2007.906629

[1]Mission details, results, and photos can be found at http://sci.esa.int and http://saturn.jpl.nasa.gov/index.cfm.

Therefore, an increased receiver input power does not necessarily improve the tracking. On the contrary, it can cause loss of symbol synchronization in the given receiver. Making the problem even worse, it is impossible to reconfigure the circuit in-orbit by telecommand or patching.

This paper illustrates a dynamical model for the symbol synchronizer, which is part of the receiver in the probe support avionics (PSA) equipment onboard the Cassini orbiter. The model describes the tracking performance of the synchronizer as a function of system-level parameters: link performance in terms of E_s/N_o; symbol transition probability P_t; and input symbol rate offset ΔF, which depends upon the relative velocity from the relay-link geometry.

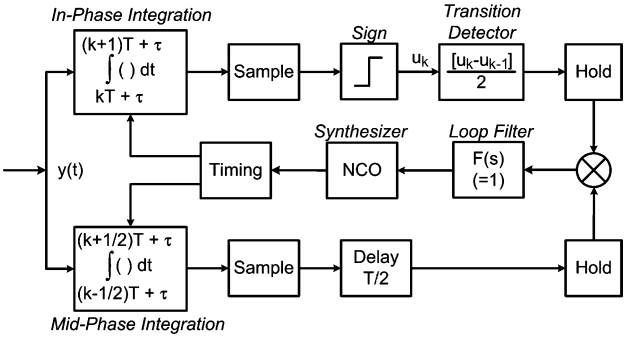

Fig. 1. ***Principle of DTTL symbol synchronizer.***

II. SYSTEM AND SIGNAL CHARACTERISTICS

The radio link [6] between the Huygens probe and the PSA onboard the Cassini orbiter employed two channels for redundancy. The link used a traditional modulation scheme with a residual carrier [5]. The symbol stream was phase-modulated onto a sinusoidal subcarrier at 131.072 kHz by applying binary phase-shift keying (BPSK). The probe nominal symbol rate was 16.384 ksymb/s, which includes the channel coding in accordance with the Consultative Committee for Space Data Systems concatenated coding scheme of Reed–Solomon (RS) (255,223) block coding and $(R = 1/2, k = 7)$ convolutional coding [7].

The receiver performs phase-coherent tracking of the residual-carrier in a second-order phase-locked loop. The phase-coherent demodulation of the subcarrier is achieved by a conventional second-order Costas loop [5]. Eventually, a traditional DTTL of first order accomplishes the symbol timing recovery and detection. The three loops for carrier tracking, subcarrier tracking, and symbol timing are closed after analog-to-digital conversion (ADC) in the digital domain of the receiver.

The PSA equipment includes a Viterbi decoder circuit, but RS decoding is not performed onboard. The received frames from the relay link, including the RS symbols, were inserted into the Cassini orbiter telemetry after being stored temporarily onboard. The RS decoding was performed on ground at the Huygens operations center.

III. SYNCHRONIZER ARCHITECTURE

The symbol synchronizer design follows the classical DTTL principle [5]. Fig. 1 shows a corresponding architectural block-diagram with the typical midphase and in-phase symbol integration paths. Whenever there is a transition between adjacent symbols, the corresponding midphase integration provides an update of the estimated offset (or misalignment) of the actual time of sampling relative to the center of the input symbol. This error signal is then used to adjust the frequency synthesizer that provides the sampling clock. For the Huygens application, a first-order DTTL is sufficient, i.e., there is no need for a loop filter $(F(s) = 1)$, provided the loop gain is set adequately. If the update to the input of the synthesizer is provided less frequently, the phase excursion or timing misalignment is increasing for a given frequency offset of the input symbol stream.

IV. ANALYTICAL MODEL

The Huygens relay link receiver has a wrong loop parameter setting in the symbol synchronizer implementation, causing too narrow a loop bandwidth. This degrades dramatically the tracking capability of the first-order loop.

For the analytical model of the symbol synchronizer, the (absolute) timing error $\epsilon(t)$ is represented by the normalized offset $\lambda = (\epsilon/T)$, where T is the symbol duration on the channel. Given that the synchronizer is a digital implementation of the DTTL, the evolution of the process λ can be represented by a first-order nonlinear difference equation

$$\lambda_{k+1} = \lambda_k - Kg\left(\lambda_k, \frac{E_s}{N_o}, P_t\right) - n(\lambda_k) + \overbrace{\left(\frac{v}{c}\right) + dF_{\text{Probe}} - dF_{\text{NCO}} - dF_{\text{TCXO}}}^{F_{\text{total}} = dF_{\text{eff}}} \quad (1)$$

with (an arbitrary) $\lambda_{k=0} = 0$; the relative velocity v from the geometry; c for the speed of light; $dF_{\text{NCO}} = -3.754$ ppm is a known DTTL preset offset of the numerically controlled oscillator (NCO); $dF_{\text{TCXO}} = -1.0$ ppm is the assumed receiver reference-clock worst case offset of the temperature-controlled crystal oscillator (TCXO). The factor K is the open-loop gain, which will be discussed in Section V. Because of the narrow loop bandwidth, any jitter effect due to additive channel noise $[n(\cdot)$ in (1)] is negligible for the analysis of the timing-error process.

Fig. 2 is a block-diagram of the system equation in (1) with $K = (4B_L T)$, where B_L is the one-sided noise

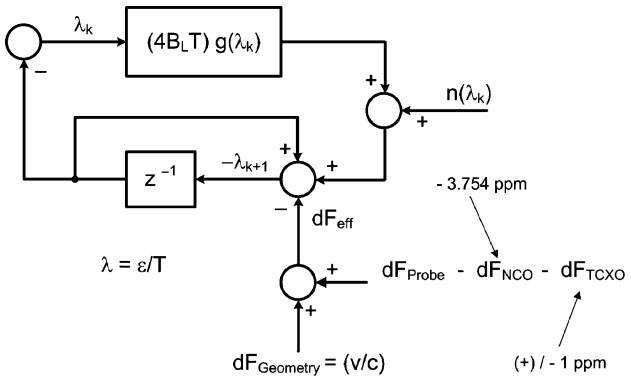

Fig. 2. ***DTTL dynamical model.***

bandwidth of the loop for a unit gain-slope of $g(\cdot)$ at $\lambda = 0$. The frequency tracking performance of the DTTL will be determined as a function of the relative offset at the DTTL input

$$\triangle F = \left(\frac{v}{c}\right) + dF_{\text{Probe}} \tag{2}$$

where (v/c) is the Doppler effect and dF_{Probe} represents potential offsets of the symbol clock generator in the transmitter onboard the Huygens probe. A relative velocity v of 300 m/s, or 1080 km/h, causes a relative input frequency offset of 1 ppm.

The function $g(\cdot)$ is the detector characteristic, or *S*-curve [5]. It represents the stochastic average of the *estimated* timing-offset $(\hat{\lambda})$ conditioned on the *actual* offset (λ). Its derivation requires statistical averaging of $\hat{\lambda}$ with respect to both noise and the random input symbol stream. The integrations in Fig. 1 are performed at symbol rate (16.384 ksymb/s) while the effective loop bandwidth related to the dynamics of the process λ is less than 1 Hz. For the time-discrete digital implementation of the DTTL, the integrate-and-dump filters process a large number (i.e., 248) of samples per each symbol. Therefore, to obtain the detector characteristic $g(\cdot)$, we apply an approximation by a time-continuous loop operation. Although in this narrow-band case the difference equation for λ_k in (1) may be considered deterministic and, thus, the influence of the loop-noise $n(\lambda_k)$ on the tracking performance is negligible, the derivation of the *S*-curve still requires statistical averaging. The result for $g(\cdot)$ in (1) is given by

$$g(\cdot) = (2P_t) \cdot \left[\lambda_k - \frac{1}{2}[1 + 2\lambda_k - P_t(1 - 2\lambda_k)] \cdot Q + \frac{1}{2}(1 - P_t)(1 - 2\lambda_k) \cdot Q|_{\lambda_k = 0}\right] \tag{3}$$

with

$$Q = \frac{1}{2} erfc\left((1 - 2\lambda_k)\sqrt{\frac{E_s}{N_o}}\right) \tag{4}$$

and $erfc(\cdot)$ for the complementary error function [8]. The $g(\cdot)$ in (3) is valid for $0 \le \lambda \le 0.5$; it needs to be periodically extended with respect to λ for the actual $g(\lambda)$ according to

$$\begin{aligned} g(\lambda) &= -g(-\lambda) \\ g(\lambda \pm n) &= g(\lambda) \text{ for integer } n. \end{aligned} \tag{5}$$

An *S*-curve $g(\cdot)$ for the DTTL is reported also in [5], which, however, does not include a detailed analysis and ignores the important dependence on P_t as a variable. References [9] and [10] provide a rigorous derivation of $g(\cdot)$. Fig. 3 shows the nonlinear *S*-curve normalized with respect to $(2P_t)$ and with E_s/N_o and P_t as parameters. The remaining dependence of the normalized $g(\cdot)/(2P_t)$ on P_t is negligible for $(E_s/N_o) > 8$ dB and, therefore, the corresponding curves overlap in Fig. 3.

V. OPEN-LOOP GAIN

The parameter K in the model (1) represents the critical open-loop gain that determines the loop bandwidth of the DTTL. It strongly depends on the implementation details and can be summarized as follows:

$$K = \sqrt{2}A \cdot \underbrace{G_{\text{ADC}}}_{\text{ADC-Gain}} \cdot \underbrace{K_{\text{SL}}}_{\text{Short-Loop}} \cdot \underbrace{[4 \cdot (F_{cl}T)]}_{\text{Pre-Det./Decim.}} \cdot \underbrace{K_{\text{SW}}}_{\text{Scaling}} \cdot \underbrace{\frac{1}{2}}_{\text{Barrel}} \cdot \underbrace{\left[\frac{(F_{cl}/8)}{2^{20}}T\right]}_{\text{NCO}} \tag{6}$$

where $\sqrt{2}A$ is the peak signal amplitude at the noncoherent AGC output. Section VII relates this amplitude to E_s/N_o. Further, it is $(1/T) = 16.384$ ksymb/s for the nominal symbol rate and $F_{cl} = 4.0625$ MHz is the clock frequency, which drives the digital symbol synchronizer. The remaining parameters in (6) are as follows.

1) ADC Gain: The ADC has a width of 4 bits. Its wideband input is dominated by noise, which is power-controlled by an analog noncoherent AGC to avoid excessive clipping effects at the ADC. The total ADC input power is 7 dBm or 0.5 V rms across 50 Ω. This implies an ADC gain [11] of $G_{\text{ADC}} = 7.59\ \text{V}^{-1}$.

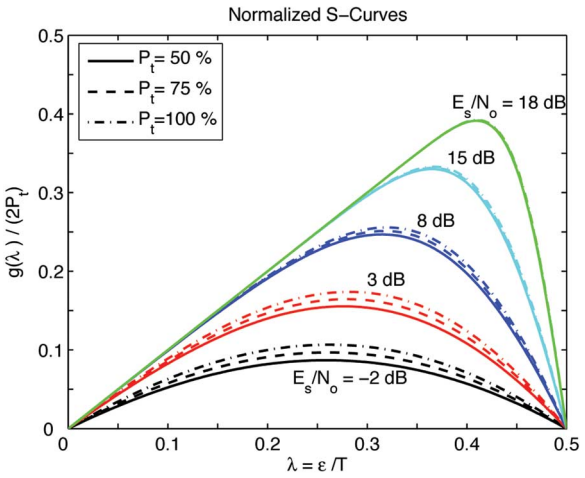

Fig. 3. ***Normalized DTTL detector characteristic with*** E_s/N_o ***as parameter.***

2) Short-Loop Closure: The intermediate-frequency (IF) signal $s_{\mathrm{IF}}(t)$ at the noncoherent AGC output is phase-modulated by the BPSK-modulated sinusoidal subcarrier. The applied modulation can be represented by

$$s_{\mathrm{IF}}(t) = \sqrt{2}A\sin[2\pi f_{\mathrm{IF}}t + ma(t)\sin(2\pi f_{\mathrm{SC}}t) + \varphi]$$

with f_{IF} and f_{SC} for the IF and subcarrier frequency, respectively; $a(t)$ for the binary (± 1) symbol stream; m for the modulation index; and φ is an arbitrary constant phase.

The receiver applies in the digital domain a short-loop closure with a corresponding conversion factor for the input signal amplitude. The parameter K_{SL} in (6) is proportional to the square root of the ratio between data power P and total signal power A^2. Therefore, K_{SL} depends upon the modulation index m with

$$K_{\mathrm{SL}} \sim C_{\mathrm{SL}}\sqrt{\frac{P}{A^2}} = C_{\mathrm{SL}}\sqrt{2J_1^2(m)} \tag{7}$$

with C_{SL} for the amplitude conversion (gain/loss) factor of the short-loop downconversion and demodulation process and $J_1(\cdot)$ for the first-kind Bessel function of first order. The conversion factor can be obtained by simulation techniques or can be approximated by $C_{\mathrm{SL}} \approx (2/\pi)\sqrt{2}$. Additionally, the parameter K_{SL} will be used in Section VI to calibrate the DTTL model for hardware implementation losses.

3) Predetection/Decimation: The hardware implementation performs a predetection by summing four digital samples from the ADC. Instead of the time-continuous integrators shown in Fig. 1, the digital DTTL circuit obtains the sum of a large number ($F_{cl}T = 248$) of samples of predetection outputs during each symbol period T.

4) Scaling: The output from the decimation, or accumulation process, is input to a microprocessor, which hosts the software for closing the loop operations of the DTTL (as well as the residual-carrier and subcarrier recoveries). For the timing recovery (first-order loop), the software performs a simple scaling by a factor of $K_{\mathrm{SW}} = 2^{-6}$. It is this selected scaling parameter that is incompatible with the original Huygens mission link geometry and the relatively high symbol rate. An additional complication is imposed by a digital coherent AGC function (Section VII), which can trigger a further reduction of the scaling factor towards 2^{-7}, or even 2^{-8}. This rescaling is implemented by a simple bit-shift operation.

5) Barrel Shifter: For the DTTL loop-closure, the microprocessor writes the scaled control signal to a register, which interfaces with the NCO in the DTTL. In order to accommodate both the results of the sign information from the in-phase path and the midphase accumulation into the same register, the control signal is effectively divided by two.

6) NCO: The NCO as part of the DTTL synthesizer in Fig. 1 is 20 bits wide and is clocked at ($F_{cl}/8$). The corresponding term in (6) represents the incremental minimum of cycles of the NCO output during the update period T of the NCO and DTTL.

VI. MODEL CALIBRATION

During the development of the receiver, the performance testing had revealed an implementation loss L_{det} of 2.75 dB, and in special cases even 3 dB, which includes the convolutional decoding. Although this loss is significant (filed in a noncompliance report), it was nevertheless accepted during the development because of comfortable link margins for the original relay-link geometry. The tests verified the system noise figure and required an increase of input signal power by 2.75–3 dB to achieve the theoretical bit error rate (after convolutional decoding). The unit tests had applied frequency uncertainties with respect to the carrier, but unfortunately not with respect to the input symbol rate; otherwise, the tracking deficiency of the symbol synchronizer would have been discovered. Therefore, in these tests, the DTTL operated in the linear region close to $\lambda \approx 0$, where the detector characteristic $g(\cdot)$ in Fig. 3 for mission-relevant E_s/N_o is almost independent of E_s/N_o.

After the detection of the anomaly and when applying similar conditions, the in-orbit tests found a consistent implementation loss L_{det} of approximately 3.1 dB.

Part of this overall loss L_{det} is due to losses in the receiver front end, in the short-loop closure—i.e., the downconversion and demodulation behind the ADC—and

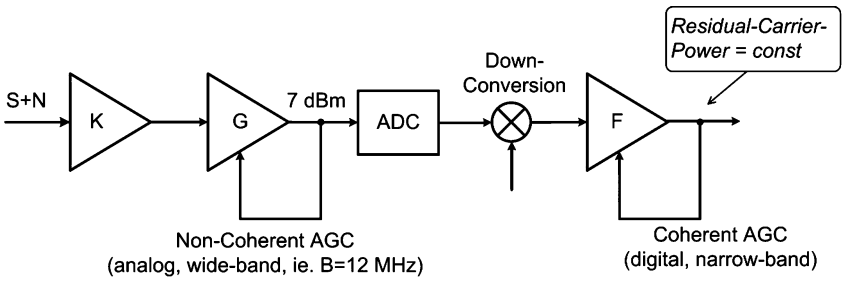

Fig. 4. ***Principle of automatic gain controls in receiver; first downconversion not shown.***

eventually in the DTTL itself. This partial loss $L_{\rm im}$ of approximately $L_{\rm im} = 1.8$ dB does not include the losses behind the DTTL, such as from the convolutional decoding. $L_{\rm im}$ is incorporated into the parameter $K_{\rm SL}$ in (6) and (7) in order to calibrate the effective open-loop gain of the DTTL model for the hardware implementation loss

$$K_{\rm SL} = C_{\rm SL} \cdot \sqrt{2J_1^2(m)} \cdot 10^{-L_{\rm im}/20}. \tag{8}$$

This a-priori knowledge of the approximate loss relevant for the DTTL performance was confirmed by the in-orbit test results reported in Section IX.

VII. AUTOMATIC GAIN CONTROL

The probe relay-link receiver includes a wide-band noncoherent analog AGC [12] that maintains a constant 7-dBm total power of signal plus noise at the ADC input. In the digital domain of the receiver, a second and narrow-band coherent AGC estimates the power of the residual carrier. The latter AGC implies at the input of the carrier recovery a signal level that is constant across the operational range of the receiver. This allows for a nearly constant effective bandwidth of the carrier-recovery phase-locked loop. Fig. 4 shows the basic concept of the two AGC functions in the receiver.

For the 50-Ω system, in (6) the peak signal amplitude $\sqrt{2}A$ (in V) at the output of the noncoherent AGC is given by

$$\sqrt{2}A = \frac{10^{P_{in}/20}}{\sqrt{10}\sqrt{1+\mathrm{SNR}^{-1}}} \tag{9}$$

with $P_{\rm in} = 7$ dBm for the total power and SNR for the (linear) signal-to-noise ratio, respectively, at the noncoherent AGC input. The SNR (in dB) is related to E_s/N_o according to

$$\mathrm{SNR} = (E_s/N_o) - 10\log(BT) - 10\log\left[2J_1^2(m)\right] \tag{10}$$

where $B = 12$ MHz is the nominal noise bandwidth at the noncoherent AGC input. This bandwidth may still be liable to vary in the in-flight equipment, since smaller bandwidths of 9.8 and 11 MHz were measured on engineering and qualification models of the receiver. Such variation would imply a slight increase in SNR.

Further to the loop-tracking performance, if the residual-carrier power at the input of the coherent AGC increases, the AGC will accordingly reduce its gain. When the AGC gain decreases below a specific preset value, the microprocessor that provides the NCO control word will reduce the scaling factor $K_{\rm SW}$ in (6) by a factor of two. If the residual-carrier power at the AGC input is increasing even further, the scaling $K_{\rm SW}$ will eventually be affected by an additional division by two.

The DTTL performances are evaluated in terms of the link parameter E_s/N_o as an independent variable. However, the switching points triggered by the coherent AGC depend upon the absolute power of the residual carrier at the coherent AGC input, which for a given E_s/N_o is a function of the modulation index m. Consequently, the switching points, if defined in terms of (E_s/N_o) values, are changing [13] with the modulation index.

For the nominal modulation index $m = 1.34$ rad, the two relevant switching points of the DTTL dwell at $(E_s/N_o) = 8.64$ and 14.84 dB, respectively. Both switching points were tested and verified during the in-orbit tests after corresponding link calibrations.

The modulation index m is subject to its own uncertainties and variations as a function of the probe temperature, for instance. An increased (decreased) index requires an increase (decrease) in E_s/N_o for a constant absolute residual-carrier power; this assumes a constant noise density N_o. Therefore, the switching points are moving to higher (lower) values and, fortunately, into the same direction as the operating point E_s/N_o. However, the absolute shifts are not identical [13].

Similarly, also a change of input noise temperature or N_o has some impact. Assuming all other parameters constant, an increase in N_o by 3 dB implies a reduction of E_s/N_o from 8.6 dB, as an example just below the first switching point, to 5.6 dB. The total power at the noncoherent AGC input is dominated by noise ((S/N) < −5 dB).

Therefore, the increase in N_o by 3 dB implies a corresponding decrease by 3 dB in gain of the noncoherent AGC. Such gain adjustment will provide a decreased residual-carrier power to the input of the coherent AGC, which subsequently will apply a 3-dB higher gain. For the coherent AGC, this represents an input operating point, which is 3 dB below the first switching point.

Maintaining this increased noise temperature and N_o, it would need an increase of signal power, assuming constant modulation index m, by the same 3 dB in order to increase the E_s/N_o towards the original 8.6 dB. As a result, the coherent AGC will operate again close to the first switching point.

This example indicates that the switching points triggered by the coherent AGC are determined by the ratio E_s/N_o and the modulation index m. Hence, the switching points do not depend any further upon the absolute value of E_s, or N_o.

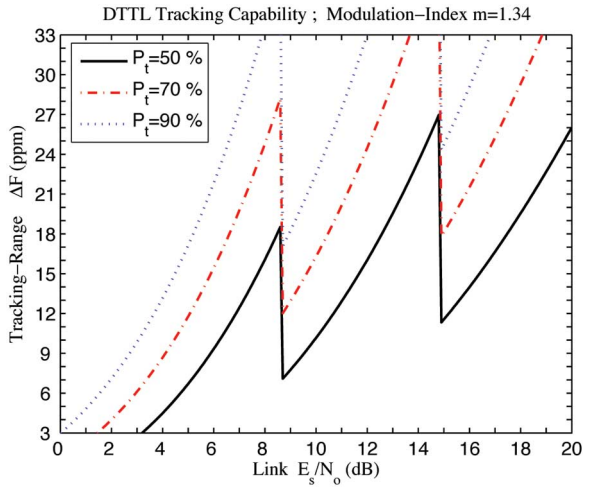

Fig. 5. ***Model for tracking range of the DTTL.***

VIII. DTTL TRACKING PERFORMANCE

The characteristics of the DTTL can be subdivided into 1) static tracking performance and 2) the dynamics of cycle-slipping for the case of the input offset exceeding the tracking range.

A. Steady-State Tracking Limits

Depending upon the system parameters of (E_s/N_o), P_t, and $\triangle F$ in (2), the process λ_k in (1) can reach a steady-state value

$$\lambda_{k+1} \rightarrow \lambda_k \rightarrow \text{const.}$$

In this case, the first-order DTTL is capable of compensating for the input frequency offset, at the penalty of a remaining static timing error, which can be very large and degrades significantly the symbol detection performance [9], [10]. In order to avoid cycle-slipping, the maximum allowed input frequency offset is given by

$$\begin{aligned}\Delta F_{\max} &= \left(\frac{\nu}{c}\right)_{\max} + dF_{\text{Probe}} \\ &= K \cdot g_{\max}\left(\frac{E_s}{N_o}, P_t\right) + dF_{\text{NCO}} + dF_{\text{TCXO}} \quad (11)\end{aligned}$$

where $g_{\max}(\cdot)$ is the maximum of the S-curve in (3) with respect to λ, for a given E_s/N_o and P_t. Note that dF_{NCO} and dF_{TCXO} are negative offsets from Section IV. The gain K depends upon E_s/N_o as per (6) and Section VII, with the specific switching points triggered by the coherent AGC. For three cases of transition densities $P_t = 50\%$, 70%, and 90%, the curves in Fig. 5 represent the tracking range in terms of maximum input frequency offset, with the discontinuities implied by the switching points at $(E_s/N_o) = 8.64$ and 14.84 dB. Similar curves were used to define several cases for the in-orbit tests and performance mappings of the DTTL in the Huygens relay-link receiver.

B. Symbol Transition Density on Channel

In Fig. 5, the tracking performance of the DTTL improves with increased transition density P_t. Random source data in the data field of the transfer frame [7] implies a P_t close to 50% on the channel after the convolutional encoding. This density can be increased by inserting source packets, which include only "0" data bits. Long sequences of zeroes are converted by the convolutional encoding into sequences with alternating zeroes and ones with a "local" transition density of 100%. (The encoder performs symbol inversion on the output path of generator polynomial G2 [7].) This increases the average P_t, and thus improves the DTTL performance. However, the insertion of "0" data packets requires an obvious reduction in effective telemetry capacity; therefore, it is not a viable solution to the DTTL problem.

C. Cycle Slipping

For parameter sets of (E_s/N_o), ΔF, and P_t that imply in Fig. 5 operating points above the corresponding "sawtooth" curve, the DTTL cannot compensate the input frequency offset and, instead, performs periodic cycle-slipping. Whenever the offset process λ passes an odd multiple of 0.5, i.e., $(2j + 1)0.5$, a symbol will be lost or skipped in the detection process.

The model in (1) allows one to simulate the time interval, in terms of channel symbols, which the synchronizer takes to complete one cycle-slip. The Huygens telemetry transfer frames [7] are 16 384 symbols long. The DTTL dynamics are much slower than the symbol rate.

Therefore, the time-discrete difference equation (1) representing the hardware implementation of the DTTL may be approximated by a first-order differential equation for a time-continuous process $\lambda(t)$. Simple integration can be applied to obtain the time T_{cs} needed by the DTTL to complete one cycle-slip during which the normalized process $\lambda(t)$ changes by one

$$T_{cs} = \frac{1}{16384} \cdot \left\{ \int_0^{0.5} \frac{d\lambda}{-K \cdot g(\lambda, (E_s/N_o)) + F_{\text{total}}} + \int_{0.5}^{1} \frac{d\lambda}{K \cdot g(1-\lambda, (E_s/N_o)) + F_{\text{total}}} \right\}$$

with T_{cs} given in terms of number of received transfer frames.

The cycle-slipping imposes loss of one symbol in the detection process per each cycle-slip. This symbol loss subsequently upsets the channel decoding process [4] and the telemetry frame synchronization in the PSA. The implied data loss in terms of affected transfer frames per cycle-slip can be related to the time interval T_{cs} [4].

The model allows one to simulate the tracking performance of the DTTL in a dynamic environment of time-varying parameters such as E_s/N_o resulting from pendulum motion of the probe below the parachute during the descent.

IX. MODEL VALIDATION BY IN-ORBIT TESTS

Fig. 6 compares the DTTL performance model with in-orbit test results. Below the limiting "sawtooth" curve, the symbol synchronizer does not show unlock phenomena ($\triangle$-points), and the data are successfully decoded. Above the "sawtooth" curve, the system cannot track the input offset ΔF ($\triangledown$-points); intermittent performances are indicated by O-points. Overall, the tests showed very good match of the model and the actual synchronizer performance. It should be noted that Fig. 6 is not any curve-fit of a model to measurement points; it is a comparison between the a-priori established model and subsequent test results. Several in-orbit test campaigns were conducted to validate the synchronizer model.

X. RECOVERY OF PROBE RELAY LINK

The design of the new probe descent relative geometry had to ensure the corresponding trajectory of input frequency offset ΔF and E_s/N_o to fall between or below the left and right "sawteeth" of the symbol synchronizer profile in order to minimize the risk of detection losses implied by entering a "sawtooth."

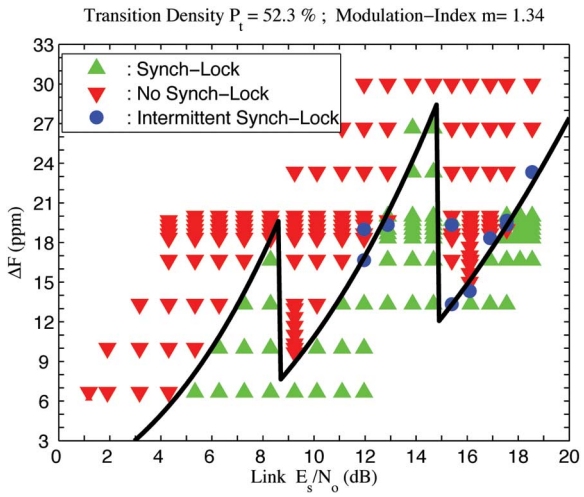

Fig. 6. *Comparison of synchronizer model and in-orbit test results.*

The concept of common link performance budgeting ensures under all mission conditions a *minimum* signal strength, which is necessary for the radio link availability and data retrieval. However, for the Huygens probe mission, the peculiar symbol synchronizer imposed the additional constraint of a *maximum* allowed E_s/N_o in order to avoid entering the "sawtooth" on the right side in Fig. 6. Especially this upper limit and the need to optimize the time-varying link conditions for the available corridor between the "sawteeth" required a radio link modeling that is conceptually different from the usual budgeting and corresponding margin policies.

A worst case example for the revised mission profile between the "sawteeth" is illustrated in Fig. 7. It had been applied during an in-orbit test. Each green arrow-marker represents the average of E_s/N_o during a 10-min time interval along the simulated probe descent of 3 h. For each time interval, the left and right crosses indicate the absolute minimum and maximum values of E_s/N_o, respectively. The test showed no data losses, except during the very last minutes inside the forbidden region.

The relay-link modeling [13] was based upon two elements.

- A static reference link budget (RLB) establishes a reference point for E_s/N_o. This budget includes the usual electrical link parameters. However, the following entries all depend upon the relative geometry and, thus, they are time-varying: 1) relative distance between Huygens probe and Cassini orbiter; 2) probe antenna gain; 3) polarization mismatch between probe antenna and Cassini high-gain antenna (HGA); and 4) HGA pointing offset. For these parameters, the RLB considers reference values, or place holders, and without variations and uncertainties' being considered. Table 1 shows

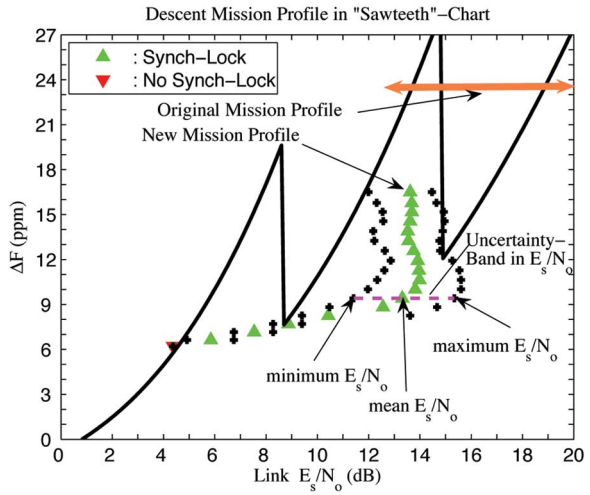

Fig. 7. *Example of the revised relay-link profile (worst case) inside the symbol synchronizer operational region.*

an RLB summary for one of the two radio channels and obtains a nominal *reference* E_s/N_o.

- A descent trajectory analysis tool (DTAT) models the variation and uncertainties of the geometry-dependent entries, for which the RLB considers reference values. Accordingly, the DTAT adjusts the reference E_s/N_o from the RLB to the real E_s/N_o predicted for the probe descent. It also models the corresponding band of estimated variation or uncertainty of E_s/N_o.

The reference E_s/N_o from the RLB is associated with an overall uncertainty due to electrical performance variations independent of the link geometry. The total uncertainty of the predicted E_s/N_o is then determined by the DTAT combining the constant uncertainty from the RLB with the variations that are geometry and time-dependent. The detailed link-budget modeling is provided in [13].

In order to avoid data losses originated by cycle-slipping in the symbol synchronizer, the link geometry has been modified radically [4], [6]. The original mission geometry would have actually maximized the Doppler shift because the probe and the Cassini orbiter were to fly almost in line with each other during the probe descent [14] through the Titan atmosphere. This implied an almost constant effective frequency offset of approximately 24 ppm indicated in Fig. 7 by the thick arrow-bar deep inside the forbidden zone of the DTTL. For the revised relay geometry, the Cassini orbiter and Huygens probe were flying side-by-side at a large distance, as shown in Fig. 8. At the periapsis, the Doppler effect was zero. Ideally, one should have had an orbiter delay time (ODT) of 1.5 h such that the estimated 3 h of probe relay link were symmetrically around periapsis and the Doppler effect is minimum. However, due to probe aspect angle (PAA) constraints, the ODT could not be reduced to less than 2.1 h. The PAA represented an important parameter for the link design because it determined the effective probe antenna gain.

XI. A SPECTACULAR DESCENT

To the immense relief of many scientists and engineers, the Huygens probe sent dramatic images of Titan and other science data captured during its descent and after the landing on January 14, 2005. The probe cameras snapped some hundreds of images, such as in Fig. 9, and successfully sent the data to Cassini. After a nominal 2-h 19-min journey through Titan's dense atmosphere, Huygens became the first space probe to land on a moon in the outer solar system.

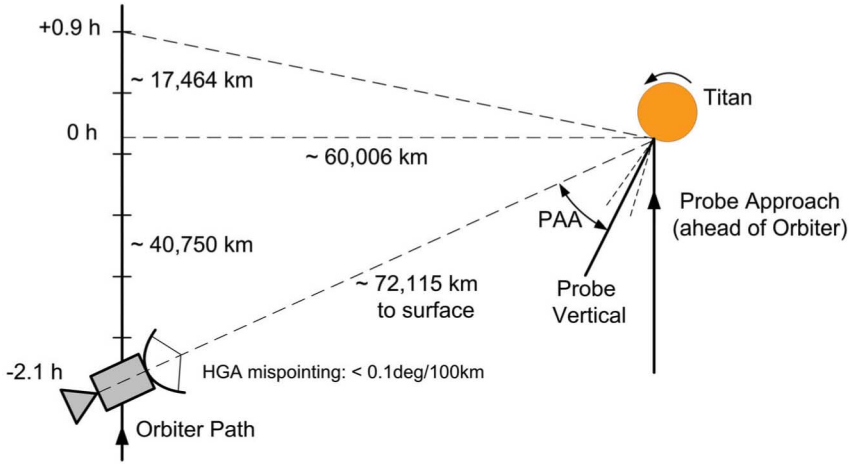

Fig. 8. *Revised relay-link geometry [6].*

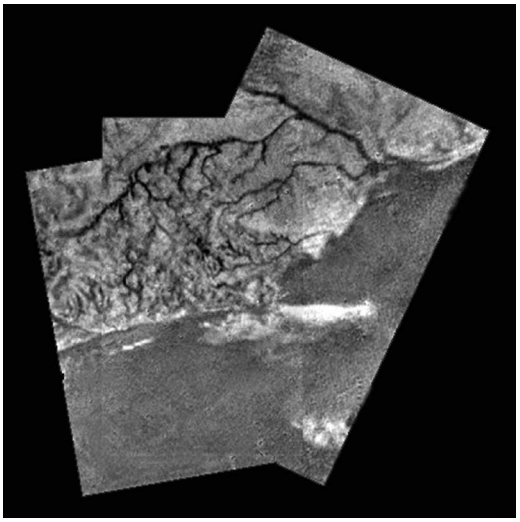

Fig. 9. *Mosaic of river channel and ridge area on Titan, taken from 16-km altitude. (Image: ESA/NASA/JPL/University of Arizona.)*

Fig. 10 shows again the forbidden operational zones (light red) where the receiver would loose lock and, consequently, the scientific probe data would have been lost inside these areas. Also included is the actual mission profile (green), which was obtained from the receiver housekeeping data during the probe mission. The match between the link predictions (envelope curves) and the actual mission profile is extremely good, which proves the accuracy of the link-budget performance predictions.

XII. LESSONS LEARNED

The incident of the Huygens relay-link implementation flaw provides some valuable lessons worth retaining [3], [4], [10] for the management of similar space missions.

The relay-link receiver is inherited from standard S-band transponders with telecommand rates not exceeding 2 kbit/s. The symbol rate of 16 ksymb/s for the Huygens application is too high for the given symbol synchronizer. During the development phase, the difference between 2 and 16 ksymb/s rates did not trigger any detailed revalidation or specific testing of the current DTTL implementation for the high symbol rate. Although an implementation may be inherited from previous projects, the existing performance values should not be taken for granted but should be verified vis-à-vis the requirements of the new application, with any adaptation to be identified for the implementation. Rigorous implementation revalidation and testing in all relevant aspects is important.

The original transponder design from which the receiver had been derived allowed for significant in-orbit reconfigurability by telecommand. This included switchable data rates and the feature to adjust the bandwidth settings for carrier as well as subcarrier loops, and indeed also for the DTTL symbol synchronizer. However, the Huygens original probe mission and operations did not require any in-orbit reconfiguration. Therefore, such receiver flexibility was considered an unnecessary design complexity and an additional risk factor avoidable for the probe mission. The reconfigurability was completely eliminated from the design by hardware settings. The

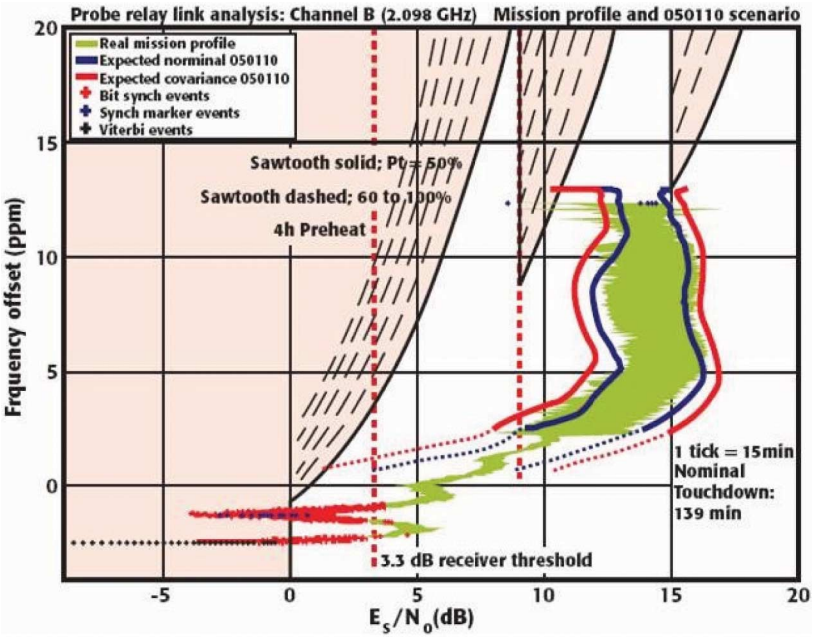

Fig. 10. *Receiver operational profile and actual mission profile retrieved from receiver housekeeping telemetry data.*

TABLE I Summary of Probe Relay Reference Link Budget (RLB)

Parameter	Value	Comment
Probe RF Power	10.66 dBW	nominal
Mismatch Loss	0.20 dB	VSWR<1.5
Circuit Losses	0.22 dB	
Probe ANT Gain	3.46 dBi	DTAT-modeled
Probe EIRP	13.70 dBW	
Frequency	2.04 GHz	
Distance	60000 km	DTAT-modeled
Free Space Path Loss	194.2 dB	
Atmospheric Loss	0.05 dB	
Tx Pol Axial-Ratio	1.62 dB	DTAT-modeled
Rx Pol Axial-Ratio	4.2 dB	max in HPBW
Polarization Loss	0.19 dB	DTAT-modeled
Propagation Loss	194.45 dB	total
Rx Ant Gain, peak	35.05 dBi	Cassini HGA
HGA Pointing Loss	0.0 dB	DTAT-modeled
Input Noise	90 K	Titan, nominal
Noise Figure	2.57 dB	System [13]
G/T	10.9 dB/K	
Rx Power	-116.15 dBm	Rx-Syst. Input
Modulation Index	1.34 rad	nominal
Data Modulation-Loss	2.50 dB	nominal
Symbol-Rate	42.14 dBHz	16.384 ksymb/s
$\mathbf{E_s/N_o}$	**14.11 dB**	Reference

receiver lacks the flexibility for modification by telecommand or to have its basic synchronization and loop parameters being patched in-orbit. Otherwise, it would have been very easy to fix the DTTL problem.

Some in-orbit reconfigurability at unit level should always be good engineering practice. This allows one to cope with unforeseen circumstances, such as in-orbit failures or the discovery of a design or implementation flaw, like in the case of the Huygens receiver, that remained undetected during the development and testing phases. The advantage of such flexibility is by far more important than an increase, if any, in design complexity. System parameters can, for instance, be stored in Flash memory and copied to registers of the hardware implementation during bootup. In future, the entire architecture of a (synchronization) system may become reconfigurable in orbit by using RAM-based field-programmable gate arrays, as an example. The configuration data are stored in Flash that can be patched by telecommand.

XIII. CONCLUSION

The symbol synchronizer in the relay-link receiver of the Huygens probe mission is incompatible with the telemetry symbol rate and the original baseline for the relay-link geometry. Extensive efforts for a verified model of the faulty implementation have been spent to cater at system level for the deficiency of the receiver. Successful verification and in-orbit testing were pursued following the detection of the anomaly. Combined with accurate radio link-budget predictions, the model defined the remaining operational corridor of the receiver.

The receiver model was the crucial key and the step forward towards the recovery of the otherwise doomed Huygens probe mission. It represented the core input to the redesign of the mission and for the radically revised orbiter/probe relative geometry. Salvaging the mission was a tremendous joint effort by many contributors from ESA, NASA/JPL, the scientific community, and industry. Eventually, this recovery allowed for the successful data retrieval from the Huygens probe during its spectacular descent to Saturn's moon Titan. ■

REFERENCES

[1] J. Oberg, "Titan calling," *IEEE Spectrum*, pp. 16–21, Oct. 2004.

[2] C. Jones and L. Popken, "Letter to editor, forum," *IEEE Spectrum*, p. 6, Feb. 2005.

[3] L. J. Deutsch, "Resolving the Cassini/Huygens relay radio anomaly," in *Proc. 2002 IEEE Aerosp. Conf.*, Big Sky, MT, Mar. 2002, vol. 3, pp. 3-1295–3-1302.

[4] "Huygens Recovery Task Force (HRTF), final report," European Space Agency, Ref. HUY-RP-12241, Issue 1, Jul. 27, 2001.

[5] W. C. Lindsey and M. K. Simon, *Telecommunication Systems Engineering*. New York: Dover, 1991.

[6] "Cassini-Huygens integrated data link report," European Space Agency, Ref. ESA-JPL-HUY-25999, Issue 2, Nov. 2003.

[7] Consultative Committee for Space Data Systems, "TM synchronization and channel coding," CCSDS 131.0-B-1, Blue Book, Sep. 2003.

[8] A. Papoulis and U. Pillai, *Probability, Random Variables, and Stochastic Processes*, 4th ed. New York: McGraw-Hill, 2002.

[9] L. Popken, "PSA Receiver Symbol-Synchronizer Dynamical Model" European Space Agency, Tech. Note, May 5, 2003.

[10] L. Popken, "Model of receiver design flaw—Crucial for Huygens space mission recovery," in *Proc. 2004 IEEE Aerosp. Conf.*, Big Sky, MT, Mar. 2004, Track 4, paper 4.1501.

[11] R. J. Baker, *CMOS, Mixed-Signal Circuit Design*. New York: IEEE Press, 2002.

[12] L. Popken and W. Kriedte, "Statistical description of non-coherent automatic gain control," *Int. J. Satellite Commun.*, vol. 11, no. 2, pp. 81–86, Mar.–Apr. 1993.

[13] L. Popken, "Probe relay RF link modeling," European Space Agency, Tech. Note, Apr. 2, 2003.

[14] L. Popken, "Rescuing Huygens," *Inst. Elect. Eng. Commun. Eng. Mag.*, vol. 4, no. 1, pp. 24–29, Feb./Mar. 2006.

ABOUT THE AUTHOR

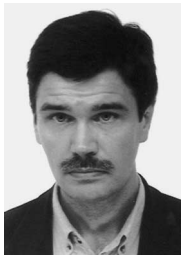

Luitjens Popken (Senior Member, IEEE) is a Senior Staff Member with the Scientific Programme Directorate, European Space Agency (ESA), Noordwijk, The Netherlands. From 1994 to 2002, he was Principal Electrical System Engineer for the avionics service module of the ESA satellite INTEGRAL (International Gamma-Ray Astrophysics Laboratory), which has been operational since 2002. In addition, he joined in 2000 as Telecommunications Expert the Huygens Recovery Task Force of ESA/NASA and later also the Huygens Implementation Team. He was responsible for the Huygens end-to-end link performance modeling and the receiver analysis. Presently, he is on a three-year secondment from ESA to the European Organisation for Astronomical Research in the Southern Hemisphere (ESO), Munich/Garching, Germany, where he is System Engineer for the Atacama Large Millimetre Array (ALMA), which is a 5600-m^2 radio interferometer array under construction in the Atacama desert of northern Chile, South America. He has authored or coauthored about 40 conference papers and journal publications. He is a Lecturer for the Erasmus Mundus SpaceMaster University Programme of the European Union.

Mr. Popken is an Associate Editor for the PROCEEDINGS OF THE IEEE Special Issue on Technical Advances in Deep-Space Communications and Tracking.