ROBUSTNESS WITH OBSERVERS

J. C. Doyle
Honeywell SRC
2600 Ridgway Rd
Minneapolis, MN, and
University of California
Berkeley, CA

G. Stein
Honeywell SRC
2600 Ridgway Rd
Minneapolis, MN, and
Massachusetts Institute of Technology
Cambridge, MA

Abstract

This paper describes an adjustment procedure for observer-based linear control systems which asymptotically achieves the same loop transfer functions (and hence the same relative stability, robustness, and disturbance rejection properties) as full-state feedback control implementations.

1. Introduction

The trouble with observers is that they tempt us, through the expedient of state reconstruction, to assign undue generality to control results proven only for the full-state feedback case. An example is the recent robustness result of Safonov and Athans [1]. This result shows that multivariable linear-quadratic optimal regulators have impressive robustness properties, including guaranteed classical gain margins of -6 db to $+\infty$ db and phase margins of ± 60 deg. in all channels. The result is only valid, however, for the full state case. If observers or Kalman filters are used in the implementation, no guaranteed robustness properties hold. In fact, a simple example has shown that legitimate LQG controller-filter combinations exist with arbitrarily small gain margins in both the positive and negative db direction [2].

In light of these observations, the robustness properties of control systems with filters or observers need to be separately evaluated for each design. Moreover, because such evaluations can come up with embarassingly small margins, a "design adjustment procedure" to improve robustness would be very desirable. The present paper provides such a procedure. We show that while the commonly suggested approach of "speeding-up" observer dynamics will not work in general, alternate procedures which drive some observer poles toward stable plant zeros and the rest toward infinity do achieve the desired objective. In effect, full-state loop-transfer properties can be recovered asymptotically if the plant is minimum phase. This occurs at the expense of noise performance.

The principal results of the paper are summarized in Section 2., where we introduce and interpret certain transfer function properties of observer-based control systems, and in Section 3, where we develop the "adjustment procedure". A simple example which illustrates these results is given in Section 4.

2. Transfer Function Properties of Observer-Based Controllers

We consider the general multivariable control loop illustrated in Figure 1. The plant is an n-th order linear system, both observable and controllable, with m inputs, p=m outputs, and no transmission zeros [3] in the right half plane. The control law consists of two transfer function matrices $H_1(s)$ and $H_2(s)$. H_2 is driven either with full-state feedback (Fig. 1A) or with an n-th order observer [4] which reconstructs the state in the usual asymptotic sense (Fig. 1B). It is clear that this overall control loop includes linear-quadratic-gaussian controllers as special cases. It also allows dynamic elements such as integrators and lag elements which may be required in more realistic control situations.

This configuration also applies to nonsquare plants for which the number of controls, m, is not equal to the number of measurements, p. For the case, $m < p$, simply augment the original control vector with (p-m) more components which are not driven by the controllers (i.e., $H_1^T = [H_{11}^T \vdots 0]$). Columns of the B matrix for these added components must, of course, be selected to introduce no unstable transmission zeros. For the case, $m > p$, select any p-dimensional subset of controls for which there are no right plane transmission zeros. Then the loop transfer properties which are established in this paper apply to this p-dimensional subset of control loops, with the remaining (m-p) loops closed.

A dashed line is shown in both Figure IA and IB in order to distinguish between elements of the loop which are part of the controller and those which are part of the plant. Since we design and implement the controller, there is relatively little uncertainty associated with it, whereas there may be significant differences between the actual plant and its model. The loop transfer functions which we examine for robustness, below, are then taken with respect to the loop breaking point, X, at the control signal interface between these two sets of elements. Very misleading robustness results can be obtained for alternate loop breaking

CH1392-0/78/0000-0001$00.75 © 1978 IEEE

points, for example Point XX. This is also shown below.

Figure 1. Linear Multivariable Control Loop

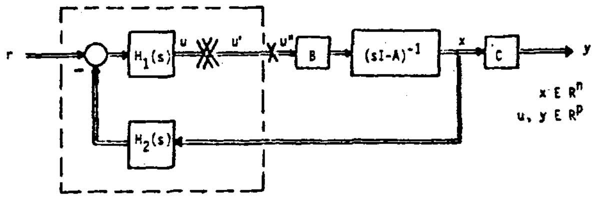

1A. Full-State Feedback Implementation

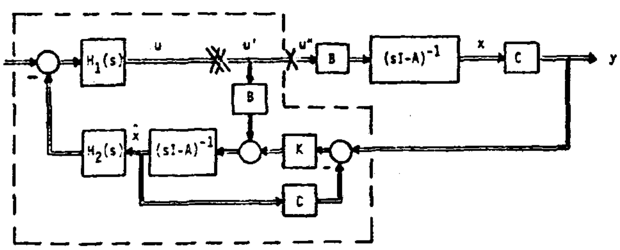

1B. Observer-Based Implementation

The following properties can be established for the above two control loop implementations:

Property 1
The closed loop transfer function matrices from command r to state x are identical in both implementations.

Property 2
The loop transfer function matrices from control signal u' to control signal u (loops broken at Point XX) are identical in both implementations.

Property 3
The loop transfer functions from control signal u" to control signal u' (loops broken at Point X) are generally different in the two implementations. They are identical if the observer dynamics satisfy

$$K[I + C\,(sI-A)^{-1}K]^{-1} = B\,[C(sI-A)^{-1}B\,]^{-1} \qquad (1)$$

The first two of these properties are very well known [5,6]. They can be easily verified by noting that the transfer functions from u' to x and from u' to $\hat{x}$ are identical because the nominal error dynamics of the observer are not controllable from u'. Hence, the error dynamics are not excited by inputs r to the closed loop system or by inputs u' to the system with loop broken at point XX.

The first two properties are also the source of much of the temptation surrounding observers, however. We see that input/output properties are the same and even certain loop transfer functions are the same. The latter promise equal relative stability properties, equal tolerance to uncertainties (robustness), and equal disturbance rejection properties. What more could we ask for? The problem, of course, is that the loop transfer properties are the same at Point XX, inside our own control implementation where only masochists would insert significant uncertain elements or disturbances. According to Property 3, equal loop transfer characteristics are not obtained at the control signal interface to the plant, Point X, where Nature gets to insert uncertainties and disturbances. It is at this point that robustness properties must be measured, and, as seen in [2], it is here that observer-based implementations can fall well short of our objectives.

The fact that loop transfer functions will in general be different at point X follows by noting that, unlike before, the observer error dynamics do get excited in response to inputs u" with loops open at X. The more interesting fact is that such differences are avoided if equation (1) holds. This latter result is apparently not as well known, so a simple derivation is given in Appendix A. It is important because it offers a way to adjust observers so that full-state loop transfer characteristics are recovered at Point X. In particular, suppose the observer gains are parameterized as a function of a scalar variable q. Let this function, K(q), be selected such that as $q \to \infty$

$$K(q) \quad \to \quad q\,BW \qquad (2)$$

for any nonsingular matrix W. Then equation (1) will be satisfied asymtotically as $q \to \infty$. The resulting observer error dynamics will have limiting poles given by roots of the polynomial

$$\psi(s) = \det(sI-A)\det\;[I + qC(sI-A)^{-1}BW]. \quad (3)$$

P of these roots will tend toward the P finite transmission zeros of the plant, i.e. the zeros of polynomial

$$\psi(s) = \det(sI-A)\det\;[C(sI-A)^{-1}B]$$

which are stable by assumption, and the rest will tend to infinity. It is clear from this that the commonly suggested approach of making all roots of the error dynamics arbitrarily faster is generally the wrong thing to do.

3. An Observer-Adjustment Procedure

Equation (2) defines the required limiting characteristics of an adjustment trajectory, K(q), which changes arbitrary initial nominal observer gains, K(0), with poor robustness properties into better gains asymptotically. We still need to define details of such trajectories.

A basic requirement for every point of an adjustment trajectory is stability of the observer error dynamics. Clearly, if we violate this requirement, overall closed loop stability is also lost. (Note that this does not mean that the net compensator within the dashed lines of Figure 1B needs to be stable). One way to assure stable error dynamics is to restrict the observer to be a Kalman filter for some set of noise parameters. That is, let

$$K(q) = \Sigma(q)\, C^T R^{-1} \tag{4}$$

with $\Sigma(q)$ defined by the Riccati equation

$$A\Sigma + \Sigma A^T + Q(q) - \Sigma C^T R^{-1} C\Sigma = 0 \tag{5}$$

As usual we take $Q = Q^T \geqslant 0$ and $R = R^T > 0$ with $(A, Q^{\frac{1}{2}})$ and (C,A) stabilizable and observable respectively. For Kalman filters, these matrices represent given process noise and measurement noise intensities. Here they are treated more freely as design parameters which we can select to suit broader purposes. In particular, let

$$Q(q) = Q_o + q^2 BVB^T \tag{6}$$

$$R = R_o \tag{7}$$

where Q_o and R_o are noise intensities appropriate for the nominal plant, and V is any positive definite symmetric matrix. With these selections, the observer gain for q = 0 corresponds to the nominal Kalman filter gain. However, as q approaches infinity, the gains are seen from (5) to satisfy,

$$K R K^T \rightarrow q^2\, BVB^T$$

and

$$K \rightarrow q\, B\, V^{\frac{1}{2}} (R^{\frac{1}{2}})^{-1} \tag{8}$$

where $V^{\frac{1}{2}}$ denotes some square root of V, i.e. $(V^{\frac{1}{2}})^T V^{\frac{1}{2}} = V$ and, similarly, $R^{\frac{1}{2}}$ is some square root of R. Since (8) is a special case of (2), it follows that the adjustment procedure defined by (4)-(7) will achieve the desired robustness-improvement objective.

Note that the second term in equation (6) can be interpreted as extra process noise added directly to the control input of the plant. Within the constraints of Kalman filter mathematics, such "fictitious noise" is a natural mechanism to represent uncertainties at this point of the control loop. It is nice to know that the resulting filter design actually responds with a corresponding robustness improvement. Note, however, that arbitrary increases of the existing noise matrix (i.e., $Q = (1 + q^2)\, Q_o$ or addition of arbitrary full rank noise process (i.e., $Q = Q_o + q^2 W$ with $W = W^T > 0$) which are often suggested as other intuitive robustness improvement methods, will not in general produce the desired effect.

Finally, we note that the use of Kalman filter equations in the adjustment procedure is not fundamental. The filters merely provide a convenient way to define a K(q) function which assures stability along the entire adjustment trajectory and has the desired limiting behavior (2). Any other procedure (pole placement, for example) with the same properties could be used as well. We emphasize, however, that both stability along the trajectory and asymptotic behavior must be achieved. Hence, such "obvious" choices as

$$K(q) \equiv q\, BW$$

will only work for special systems which are stabilizable with high gain output feedback alone. The Kalman filter choice (8), in contrast, works for all controllable, observable minimum-phase plants.

4. An Example

To illustrate the observer properties and adjustment procedure above, consider the following example:

Plant:

$$\frac{dx}{dt} = \begin{bmatrix} 0 & 1 \\ -3 & -4 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u + \begin{bmatrix} 35 \\ -61 \end{bmatrix} \xi \tag{9}$$

$$y = [2 \quad 1]\, x + \eta \tag{10}$$

with $E(\xi) = E(\eta) = 0$;

$$E[\xi(t)\xi(\tau)] = E[\eta(t)\eta(\tau)] = \delta(t-\tau)$$

Controller:

$$u = [-50 \quad -10]\, \hat{x} + [50]\, r \tag{11}$$

The plant in this example is a (harmless) stable system with transfer function.

$$\frac{y}{u}(s) = \frac{s + 2}{(s + 1)(s + 3)} \tag{12}$$

The contoller happens to be a linear-quadratic one, corresponding to the performance index

$$J = \int_0^\infty (x^T H^T H x + u^2)\, dt \tag{13}$$

with

$$H = 4\sqrt{5} \quad [\sqrt{35} \quad 1]$$

It places the closed loop regulator poles at

$$s = -7.0 \pm j2.0$$

A Nyquist diagram (polar plot of the loop transfer function at Point X) for the full-state design is given in Figure 2. Gain margin is infinite in both directions and there is over 85° phase margin. The design is then implemented using a Kalman filter for the given noise parameters. The Nyquist plot for the resulting observer-based controller is also shown in Figure 2. Oops... less than 15° phase margin.

In an effort to improve this margin, one adjustment to the filter that could be made is to speed it up. So, we can try moving the filter/observer poles to the left in a second-order Butterworth pattern. For the filter/observer poles at $-22 \pm 17.86j$ one gets the third Nyquist plot in Figure 2. As can be seen, the results are less than satisfactory. Not only are the margins disappearing (now less than 10 degrees) but the loop bandwidth has increased (crossover has gone from approximately 12. to 40. rad/sec).

Unless we're trying to design an explosive device, this is clearly undesirable. It gets worse as the filter gets faster. In fact, it can be shown that the margins go asymptotically to zero for large gains, while the loop bandwidth goes to infinity. The present example is not a pathological one, either. Similarly undesirable characteristics for fast filters are obtained with most systems.

When the observer adjustment procedure of Section 3 is applied to the same example, much more pleasing behavior is obtained. Following (6)-(7), we let the process noise covariance matrix be

$$Q = \begin{pmatrix} 35 \\ -61 \end{pmatrix} (35 \quad -61) + q^2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} (0 \quad 1) \tag{14}$$

We then increase q from zero until a reasonable compromise between noise performance and robustness is achieved. Some results of this process are summarized in Figure 3 and Table 1. Figure 3 shows Nyquist diagrams for $q^2 = 100$, 500, 1000, and 10,000. Margins improved with essentially no change in bandwidth as the modified loop transfer function tends toward full state optimal. Noise performance is summarized in Table 1 for the same set of q values. As expected, the error covariance of the adjusted filter with respect to the original noise increases markedly with q. However, there was not the same deterioration in state covariance.

Table 1 also documents other parameters associated with these design points - - poles of the error dynamics, margins, and filter gains. Note in particular that the filter poles tend toward the plant zero and toward infinity, as required by (3).

This adjustment procedure was also successfully applied to reconstruction of measured outputs after sensor failures for the A7-D aircraft. [8]. In this application the optimum Kalman filter produced an unstable system when tested in hybrid simulation over the A7-D flight envelope. After attempts with "ad hoc" ficititious noise adjustment procedures failed the method discussed here successfully stabilized the system. Also, the resulting error covariance properties remained closed to the optimum values.

5. Conclusions

This paper illustrates some of the difficulties one can get into by relying on observers for state reconstruction. We have concentrated on robustness properties. In general, these will be poorer for observer-based implementations than for full-state implementations. For minimum-phase systems, however, full-state robustness can be recovered asymptotically provided it is done correctly. Fast observers are not in general correct. A "fictitious noise" adjustment procedure was suggested which is.

The apparent practical value of this procedure is that it gives a simple way of trading off between noise rejection and margin recovery. When q = 0, the noise filter will be optimal with respect to the "true" (as modelled) system noise. As q increases the filter will do a poorer job of noise rejection but the closed-loop stability margins will improve. Hopefully, a satisfactory compromise can be found through the adjustment of the single parameter q. We stress that margin recovery occurs at Point X in Figure 1 -- at the control signal interface to the outside world. Asymptotically, the full-state and observer-based implementations will have the same tolerance to disturbances and uncertain elements inserted at this point. While Point X is clearly a physically important one (more inportant that Point XX, certainly), engineers who may wish to test robustness at still other points in the control loop should recognize that the recovery results may not be applicable there.

The suggested adjustment procedure is essentially the dual of a sensitivity recovery method suggested by Kwakemaak [7]. The latter provides a method for selecting the weights in the quadratic performance index so that full-state sensitivity properties are achieved asymptotically as the control weight goes to zero. In this case, however, closed loop plant poles instead of observer poles are driven to the system zeros, which can result in unacceptable closed loop transfer funcion matrices for the final system.

ACKNOWLEDGEMENTS

We would like to thank the Math Lab Group, Laboratory for Computer Sciences, MIT for use of their invaluable tool, MACSYMA, a large symbolic manipulation language. The Math Lab Group is supported by NASA under grant NSG 1323 and DOE under contract #E(11-1)-3070.

APPENDIX A: Derivation of Property 4

Referring to Figure 1A, the loop transfer function from u" to u' of the full state implementation is obtained from the relationships

$$x = \Phi (Bu'' + Fv) \tag{A.1}$$

$$u' = -H_1 H_2 x, \tag{A.2}$$

where

$$\Phi = (sI - A)^{-1} \tag{A.3}$$

$$v = -G_1G_2x \ . \quad (A.4)$$

The variables v above are not shown in Figure 1 for the sake of simplicity. They denote the (m - p) control components for which loops are not broken in the event that p < m. Matrices F, G_1, and G_2 are the control input matrix and the feedback compensator matrices for these components, respectively. If the original plant is square or can be made square by augmenting (p - m) additional control variables, then v, F G_1 and G_2 are zero identically. For either situation, (A.1) - (A.4) define the following full-state loop transfer function:

$$u' = -H_1H_2(I + \Phi FG_1G_2)^{-1} \Phi Bu'' \quad (A.5)$$

The corresponding relationships for observer-based implementations are (Fig. 1B).

$$\hat{x} = (\Phi^{-1} + KC)^{-1} \{Bu' + Fv + KC\Phi(Bu'' + Fv)\}$$
$$= (\Phi^{-1} + KC)^{-1} \{Bu' + KC\Phi Bu'' + (\Phi^{-1} + KC)\Phi Fv\}$$
$$= (\Phi^{-1} + KC)^{-1} \{Bu' + KC\Phi Bu''\} + \Phi Fv \quad (A.6)$$

with

$$u' = -H_1H_2\hat{x}$$
$$v = -G_1G_2\hat{x} \quad (A.7)$$

This gives

$$u' = -H_1H_2(I + \Phi FG_1G_2)^{-1} (\Phi^{-1} + KC)^{-1} \{Bu' + KC\Phi Bu''\} \quad (A.8)$$

Now applying the Matrix inversion lemma [9] to the $(\Phi^{-1} + KC)^{-1}$ term in this expression gives

$$u' = -H_1H_2(I + \Phi FG_1G_2)^{-1} [\Phi - \Phi K (I + C\Phi K)^{-1} C\Phi] \{Bu' + KC\Phi Bu''\}$$
$$= -H_1H_2(I + \Phi FG_1G_2)^{-1} \Phi [B - K(I + C\Phi K)^{-1} C\Phi B]u'$$
$$-H_1H_2(I + \Phi FG_1G_2)^{-1} \Phi K (I + C\Phi K)^{-1} C\Phi Bu''. \quad (A.9)$$

From (A.9) it follows that if (1) is satisfied, then the u' term on the right hand side vanishes and the u" term is identical to (A.5). Since u" is arbitrary, this establishes the claimed equality of loop transfer functions.

REFERENCES

1. Safonov, M. G., and M. Athans, "Gains and Phase Margin of Multiloop LOG Regulators, "IEEE Trans. Auto. Contol, April 1977.

2. Doyle, J. C., "Guaranteed Margins for LOG Regulators," IEEE Trans. Auto. Control, August 1978.

3. MacFarlane, A. G. J. and Karcanias, N., "Poles and Zeros of Linear Multivariable Systems: A Survey of Algebraic, Geometric, and Complex Variable Theory," Int. J. Control, July 1976, pp. 33-74.

4. Schweppe, F. C. Uncertain Dynamic Systems, Prentice-Hall, 1973.

5. Kwakernaak, H. and Sivan, R., Linear Optimal Control Systems, Wiley-Interscience, 1972.

6. Anderson, B. D. O. and Moore, J. B., Linear Optimal Contol, Prentice-Hall, 1971.

7. Kwakernaak, H., "Optimal Low-Sensitivity Linear Feedback Systems," Automatica, Vol. 5, No. 3, May 1969, p.279.

8. Cunningham, T.B., Doyle, J. C., and Shaner, D. A., "State Reconstruction for Flight Control Reversion Modes", 1977 IEEE Conference on Decision and Control, New Orleans, December 1977.

9. Householder, A. S., "Principles of Numerical Analysis," McGraw-Hill, New York, 1953.

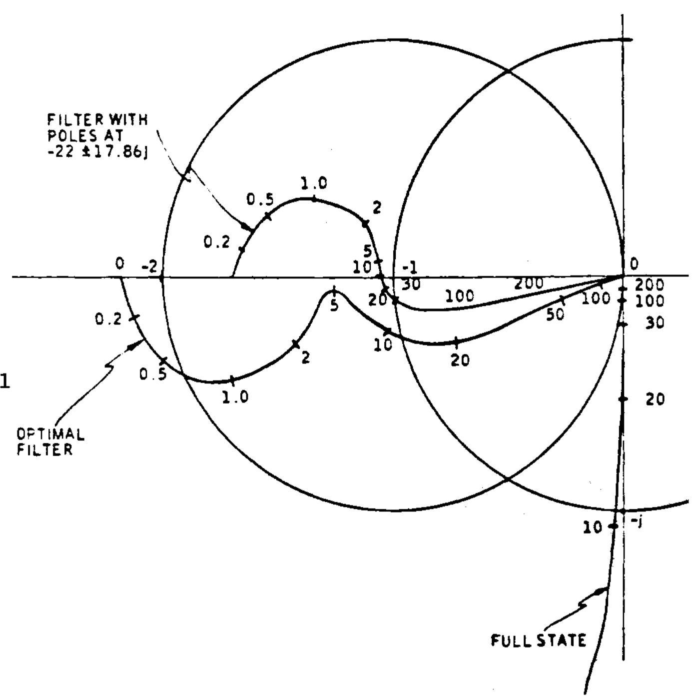

Figure 2. Loop Transfer Functions of Example: "Fast Filter" Adjustment Procedure

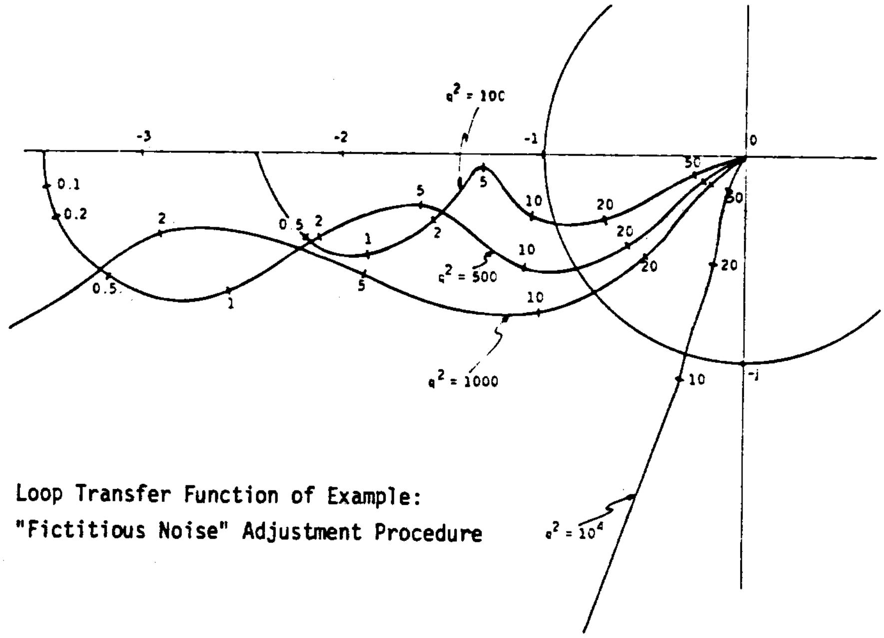

Figure 3. Loop Transfer Function of Example: "Fictitious Noise" Adjustment Procedure

	FILTER POLES	GAIN MARGIN db	PHASE MARGIN deg	ERROR COVARIANCE $E\left[(x-\hat{x})(x-\hat{x})^T\right]$	STATE COVARIANCE $E(xx^T)$	FILTER GAIN
Optimal LQG Design	$-7 \pm 2j$	- 6.75	15	97 - 163 - 163 277	221 - 613 -613 2070	30 - 50
Fast Filter Adjustment Procedure.	$-22 \pm 17.86j$	- .98	<10	6284 -12224 -12224 23788	130 - 613 -613 8517	720 -1400
Fictitious Adjustment Procedure $q^2 = 100$	-4.3 -13.1	- 7.73	19	107 - 184 - 184 319	236 -613 -613 1812	26.8 - 40.2
$q^2 = 500$	-2.9 -24	-10.9	33	163 - 301 -301 564	268 -613 -613 1497	20.4 - 17.7
$q^2 = 10^3$	-2.5 -33	13.9	42	204 -385 -385 743	285 -613 -613 1360	16.7 - 1.9
$q^2 = 10^4$	-2.1 -100	-37	74	290 -570 -570 1169	317 -613 -613 1198	6.9 84.6

TABLE 1. SUMMARY OF EXAMPLE