MECHANIZING DYNAMIC SECURITY ANALYSIS

by

Richard J. Marceau, Eng.

B.Eng. (McGill University, Montréal, Canada)
M.Sc.A. (École Polytechnique de Montréal, Montréal, Canada)

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

Department of Electrical Engineering
McGill University
Montréal, Canada

ISBN 0-315-91717-2

To my parents
Bruno and Irène Marceau

and to my wife and children
Denyse Vézina, Aude and Hugues Marceau

without whose love and support this work
could not have been done.

MECHANIZING DYNAMIC SECURITY ANALYSIS

ABSTRACT

The object of software *frameworks* is to mechanize human processes in order to accomplish high–level tasks that call upon diverse software tools. This thesis describes the ELISA framework prototype which performs power–system dynamic security analysis in the operations planning environment. ELISA mechanizes routines traditionally carried out by experts that are essential to power–system dynamic security analysis, greatly accelerating the realization of complex processes. Typically, ELISA executes appropriate load–flow and transient–stability simulations (i.e. using commercially available simulation software), performs result analysis, identifies and executes changes to the input and repeats this process until a user–defined goal, such as finding transient stability transfer limits, has been achieved.

A taxonomy of dynamic security analysis in operations planning is proposed employing the semantic net, class–object–property and rule paradigms. All of these are required to cover the full spectrum of knowledge found in the high–level goals, the process details, the complex conditional structures and the acceptance criteria which characterize dynamic security analysis. This taxonomy also describes the language of operations planners, defining not only the features presently supported by ELISA, but also providing a roadmap to future enhancements. Typical sensitivity studies are presented using a 700–bus production model of the Hydro–Québec network to illustrate the considerable leverage afforded from using ELISA–like software.

In addition, the thesis addresses the issue of how such tools can assist in performing research to improve our understanding of fundamental power systems behaviour. Using the ELISA prototype as a laboratory test bed, it is shown that the signal energy E of a network's transient response acts as a barometer to define the relative severity of any normal contingency with respect to power generation or transfer P. For a given contingency, as P is varied and the network approaches instability, signal energy increases smoothly and predictably towards an asymptote which defines the network's stability limit: This limit, in turn, permits us to compare the severity of different contingencies. This behaviour can be explained in terms of the effect of increasing power on the damping

component of dominant poles, and a simple function is derived which estimates network stability limits with surprising accuracy from two or three stable simulations.

As a corollary to this, it is also shown that a network's transient response can be screened for instability using a simple frequency–domain criterion. Essentially, this criterion requires performing the Fourier transform of a network's transient voltage response at various monitoring locations: When P is varied and the network goes beyond its stability limit, the angle of the Fourier transform's polar plot fundamentally changes its behaviour, passing from a clockwise to a counterclockwise rotational behaviour about the origin. This is confirmed by results obtained from performing stability–limit searches on the Hydro–Québec system. Used in conjunction with signal energy analysis for determining stability limit proximity, this criterion can be quite useful for mechanized security–limit–determination tools such as ELISA.

Signal energy limit estimation and the proposed stability criterion are shown to be applicable to all normal contingencies and these results hold notwithstanding the presence of many active, nonlinear elements in the network.

Keywords. Frameworks; Dynamic security; Transient stability; Stability theory

LA MÉCHANISATION DE L'ANALYSE DE LA SÉCURITÉ DYNAMIQUE

RÉSUMÉ

Les *frameworks* visent la méchanisation des processus humains et ce, en vue de faciliter la réalisation de tâches de haut niveau faisant appel à l'utilisation de nombreux logiciels. La présente thèse décrit le prototype ELISA, lequel mécanise des processus experts en matière d'analyse de sécurité dynamique des réseaux en planification de l'exploitation. Typiquement, ELISA effectue des simulations autonomes d'écoulement de puissance et de stabilité transitoire (i.e. utilisant des logiciels de simulation disponibles commercialement), analyse les résultats, modifie les données à l'entrée et répète le processus jusqu'à ce que l'objectif défini par l'utilisateur, telle la recherche d'une limite de transit, est atteint. La mécanisation de tels processus a pour effet d'accélérer considérablement leur réalisation.

Une condition préalable à la réalisation d'un tel prototype est de bien décrire le caractéristiques de l'analyse de la sécurité dynamique. Dans cette thèse, on propose une taxonomie de l'analyse de la sécurité dynamique fondée sur différents paradigmes de description des connaissances, à savoir les réseaux sémantiques, les classes et objets, ainsi que les règles. Une intégration de ces paradigmes est requise aux fins d'exprimer toute la richesse du domaine et ce, par la description des objectifs, des processus, des structures décisionelles complexes et des critères d'acceptation qui lui sont propres. De par sa nature, cette taxonomie définit également le langage des planificateurs de l'exploitation: plus particulièrement, il décrit à la fois les caractéristiques du prototype actuel et les améliorations futures qu'on pourra lui apporter. On y présente quelques exemples d'études de sensibilité sur le réseau complet d'Hydro–Québec afin d'illustrer l'important effet de levier que procure l'utilisation d'ELISA dans le milieu de la planification de l'exploitation.

De plus, la présente thèse explore comment de tels outils peuvent assister dans la poursuite de recherches visant à approfondir nos connaissances fondamentales en matière de réseaux. Utilisant ELISA comme un banc d'essai, on démontre que l'énergie du signal E de la réponse transitoire d'un réseau est un baromètre qui quantifie la sévérité relative d'une contingence vis–à–vis d'une production ou d'un transit P. Pour une contingence donnée, en variant P de

manière à s'approcher de la limite de stabilité, l'énergie du signal tend asymptotiquement vers la limite de la stabilité du réseau. Cette limite, à son tour, permet de comparer la sévérité de différentes contingences. Ce comportement s'explique en fonction de l'incidence de l'augmentation de la puissance transitée sur la composante réelle des pôles dominants. À partir de ce principe, on obtient une formule simple qui évalue la limite de stabilité d'un réseau au moyen de deux ou trois simulations stables et la valeur estimée est d'une précision surprenante.

De plus, on y propose un critère simple pour l'analyse de la stabilité transitoire. Il s'agit d'un critère appliqué au domaine de la fréquence: essentiellement, on examine la transformée de Fourier de la réponse transitoire de la tension à différents endroits dans le réseau. Quand on varie P de façon à ce que le réseau dépasse sa limite de stabilité, l'angle du graphique polaire de la transformée de Fourier modifie fondamentalement son comportement, passant d'une rotation dans le sens de l'horloge à une rotation dans le sens contraire de l'horloge. On confirme ceci par de nombreux résultats obtenus de simulations du réseau d'Hydro–Québec.

Enfin, on démontre que l'estimation des limites au moyen de l'énergie du signal et le critère de stabilité proposé plus haut sont applicables à toutes les contingences normales et que ces résultats sont valables nonobstant la présence de nombreux éléments actifs et non–linéaires dans le réseau.

Mots clés. Frameworks; Sécurité dynamique; Stabilité transitoire; Théorie de la stabilité.

ACKNOWLEDGEMENT

Words do not always come easily to engineers, especially words dealing with feelings rather than facts or observations. I hope that these few words will do justice to those who have made this thesis possible.

I first wish to express my profound gratitude to Professor F. D. Galiana, my thesis supervisor, for his wisdom, his guidance, his ready availability, his unflagging support, his gentle prodding, his generosity beyond compare and for the warm friendship that characterized our relationship. I also wish to thank Mr. Donald T. McGillis, Dr. A. Malowany, Dr. B.T. Ooi and Dr. Zames for their valuable comments and their constant encouragement.

I wish to thank Richard Mailhot of Hydro–Québec for the knowledge and insight so willingly shared of operations planning. His expertise greatly contributed to the development of the ELISA framework prototype.

I gratefully acknowledge the contributions of Florent Dénommé and Peter Czech for having initiated a fruitful collaboration between McGill University and Hydro–Québec. I hope that this marks the beginning of a long relationship.

Thanks are due to many other Hydro–Québec colleagues who accepted me in the operations planning group. Notwithstanding their heavy work load, they readily shared their experiences, answered questions, and tested or commented some new ELISA feature or idea. In particular, I would like to mention François Lévesque, Marc Langevin, Daniel Lefebvre, André Lapointe, Laval Riverin, George Blais, Julien Gagnon and Réjean St–Hilaire. Thanks are also due to my fellow Ph.D. students, in particular Lester Loud, Hossein Jovidi and Katia Campos de Almeida for their warm companionship.

Special thanks are due to the SUN network administrator at Hydro–Québec, Arline Bourdages, who was always ready to lend a helping hand. Many thanks are also due to her assistant, Christian Legault, who did such a good job on the ELISA windows user–interface.

Last but not least, I wish to thank the Hydro–Québec Bursaries Committee, and to José Salgado, former acting Director of R&D Planning, for permitting me to pursue a doctoral program. It is a privilege to pause midway through one's life to realize a lifelong ambition and I will always be grateful.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ABBREVIATION	MEANING
AC	Alternating Current
CCW	Counterclockwise
CW	Clockwise
DC	Direct Current
EHV	Extra High Voltage
EMS	Energy Management System
ELISA	Estimateur de LImites de Sécurité Automatisé
FACTS	Flexible Alternating Current Transmission System
FVP	Flat Voltage Profile
FFT	Fast Fourier Transform
FT	Fourier Transform
HVDC	High Voltage Direct Current
IEEE	Institute of Electrical and Electronics Engineers
MVA	MegaVolt–Amperes
MVAR	MegaVolt–Amperes Reactive
MW	MegaWatts
NERC	North American Electric Reliability Council
NPCC	Northeast Power Coordinating Council
OPF	Optimal Power Flow
p.u.	Per Unit
PV	Power–Voltage (i.e. at a PV bus, power and voltage are specified in the power flow equations)
PQ	Power–Reactive Power (i.e. at a PQ bus, power and reactive power are specified in the power flow equations)
RP600	Hydro–Québec power flow program
SC	Synchronous Condenser
SISO	Single Input Single Output
SLG	Single Line to Ground
SPS	Special Protection System
ST600	Hydro–Québec transient stability program
STPP	Post processor for obtaining graphical output of ST600 results
SVC	Static VAR Compensator
TOV	Temporary Overvoltage

LIST OF SYMBOLS

SYMBOL	MEANING
alphap	Voltage exponent for modelling different types of load.
B_{ij}	A constant in equation (5−14) for $H_{ij}(\omega, P)$
B_{unit}	Charging of a single generator or transformer unit
$B_{dispatch}$	Charging of a set of multiple generators or transformers in parallel
C_k	Normal contingency k
C_{0ij}	A constant in equation (5−26) for $E_{ij}(P_i)$
C_{1ij}	A constant in equation (5−23) for $E_{ij}(P_i)$
C_{2ij}	The power transfer limit in equations (5−23) and (5−26) for $E_{ij}(P_i)$
C_{3ij}	A constant in equation (5−23) for $E_{ij}(P_i)$
$\delta(t)$	An ideal or true impulse
$\delta_a(t)$	An approximate impulse
$d(t, P)$	Time− and power−dependent topology−changing disturbance function
$d_{ij}(t, P)$	Time− and power−dependent topology−changing disturbance function, for a fault at location i and monitored at location j
$D_{ij}(\omega, P)$	Fourier transform of a topology−changing disturbance function dependent on time and power, for a contingency occurring at location i and monitored at location j
$E, E(P)$	Signal energy or power dependent signal energy
E_{ij}	Signal energy for a contingency occurring at location i and monitored at location j
$E_{ij}(P_i)$	Signal energy for a contingency occurring at location i and monitored at location j, dependent on power
E_{ij}'	Per unit signal energy for a contingency occurring at location i and monitored at location j
$E_{ij}'(P_i')$	Per unit signal energy for a contingency occurring at location i and monitored at location j, dependent on power in per unit of the stability limit
$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'}$	Change in per unit signal energy between two points as a function of per unit power
$\frac{dE_{ij}'(P_i')}{dP_i'}$	Per unit signal energy slope as a function of per unit power

ε	Time length of pulse of height ε
$\mathcal{F}[\ ]$	Fourier transform
h	Impulse response
$h(t, P)$	Time– and power–dependent impulse response
$h_a(t, P)$	Approximate time– and power–dependent impulse response
$h_{ij}(t, P)$	Time– and power–dependent impulse response for a contingency occurring at location i and monitored at location j
$H(\omega, P)$	Fourier transform of a time– and power–dependent impulse response (i.e. transfer function)
$H_{ij}(\omega, P)$	Fourier transform of a time– and power–dependent impulse response for a contingency occurring at location i and monitored at location j (i.e. transfer function)
$H_{ij}(\omega, P)'$	Fourier transform of the transient response (including disturbance function) for an true impulse at the input
H_{unit}	Inertia constant of a single generating unit
$H_{dispatch}$	Inertia constant of multiple generators in a power station
I_f	Fault current
K	Area under approximate impulse (equation 5–5)
k_{ij}	A constant in equation (5–22) for σ_{ij}
K_{ij}	A constant in equation (5–16) for $R_{ij}(\omega, P)$
N	Number of dispatched generating units
Nb	Base network
Nd_i	Network degraded by a number of line sections or lines
Nd_{ij}	Network degraded by a number of voltage support elements or tie lines
p_{ij}	Dominant poles of network transfer function for a contingency occurring at location i and monitored at location j
P	Power generation or transfer
P_i	Power generation or transfer parameter which causes the power system to approach the stability limit for a contingency occurring at location i
P_i'	Power in per unit of the stability limit
P_0	Value for power at a constant power bus obtained from the power flow program
P_{unit}	Power dispatching of a single generating unit
$P_{dispatch}$	Power dispatching of multiple generators in a power

	station
r	Transient response of a power system
$r(t, P)$	Time– and power–dependent transient response of a power system
$r_{ij}(t, P)$	Time– and power–dependent transient response of a power system for a contingency occurring at location i and monitored at location j
$R(\omega, P)$	Fourier transform of the time– and power–dependent transient response of a power system
$R_{ij}(\omega, P)$	Fourier transform of the time– and power–dependent transient response of a power system for a contingency occurring at location i and monitored at location j
σ_{ij}	Dominant pole damping component (i.e. real part of p_{ij}) for a contingency occurring at location i and monitored at location j
$S(P)$	Impulse response signal energy
t_f	Fault duration time
t_s	Simulation duration time
v, $v(t, P)$	Time– and power–dependent voltage waveform
$v_{ij}(t, P)$	Time– and power–dependent voltage waveform for a contingency occurring at location i and monitored at location j
V_{st-st}	Steady–state post–contingency voltage
V	Pre–contingency voltage at point of fault; also voltage at any bus
V_f	Voltage accross the combined negative and zero sequence impedances of the network during a single–line–to–ground fault
V_1	Voltage accross the network's positive sequence impedance during a single–line–to–ground fault
V_{fd}	Desired voltage accross the combined negative and zero sequence impedances during a single–line–to–ground fault
V_{JB}	James Bay voltage phasor
V_{CF}	Churchill Falls voltage phasor
V_{Duv}	Duvernay voltage phasor
ω	Angular frequency
ω_{ij}	Dominant pole angular frequency for a contingency occurring at location i and monitored at location j

Xd'	Transient reactance of a single generator
Xd'_{eq}	Equivalent transient reactance of multiple generators
Z_0	Network zero sequence impedance as seen from the fault bus
Z_1	Network positive sequence impedance as seen from the fault bus
Z_1	Network negative sequence impedance as seen from the fault bus
Z_{unit}	Impedance of a single generator unit
$Z_{dispatch}$	Equivalent impedance of a set of multiple generators operating in parallel
Z_f	Fault impedance inserted in series with positive sequence network (i.e. the sum of the zero and negative sequence impedances) for simulating single–line–to–ground (SLG) faults; also test fault impedance
Z_{fd}	Desired single–line–to–ground fault impedance

CHAPTER 1

INTRODUCTION

1.1 SECURITY

Webster's Dictionary [Webster 1987] defines security alternatively as "freedom from danger" or "measures taken to guard against ... attack". For millennia, "humanity's struggle to survive and find security in the midst of endless change and movement" [Cleary 1989] has been a concern of practical philosophers, and the pursuit of security continues to preoccupy us in such different and unrelated areas as government, business, finance and the environment. In a world governed by change, the main obstacle to achieving security is uncertainty, a concept which embraces such disparate elements as incomplete knowledge of the state of the world, unfamiliar tasks, inaccurate predictions or imperfect agents [Talukdar & Cardoso 1988].

To guard against uncertainty, one attempts to minimize it by extrapolating from past experience, obtaining greater knowledge of the present, and attempting to exercise control over one's environment [Sun–Tzu 1993]. Successful strategists have been found to combine these with innovation, intuition and creativity [Carr 1992, Pasquier 1992] and, from this perspective, security has all of the attributes of an art. Nevertheless, security must also be regarded as a science in that measurable quantities are often critical in guiding the decision process: such quantities are useful in monitoring events, analyzing trends, and establishing conditions or thresholds for renewed action [adapted from Cleary 1989]. In a broad sense, security can thus be regarded as the art and science of ensuring survival, and strategy is the discipline which concerns the ways and means of finding security.

Because survival is an instinct of all living things, "security" is a term almost intuitively understood, particularly when applied to other living things. Beyond this context, the meaning of the word "security" is more difficult to grasp, particularly when considering very large and complex technologies such as power systems. When electric utility engineers discuss such concepts as "power systems security", "steady–state security" or "dynamic security", what are they talking about? And why is security so important to power systems?

To begin, an operating power system must also deal with uncertainty. This usually takes the form of a variety of planned or unexpected and sometimes undesirable events or disturbances. These can arise from changes in load, equipment malfunction, climactic or meteorological conditions, natural or industrial pollution, even natural catastrophes such as earthquakes, volcanic eruptions, and many others. Such occurrences perturb the normal operation of power systems, leading to electrical faults, unexpected changes in network topology, equipment or line overloads, the damage or loss of equipment, control system malfunctions, even large–scale failure leading to system shutdowns [Ewart 1978]. When reacting to unforeseen events, a power system is vulnerable to large–scale failure in two very different ways. On the one hand, a disturbance may lead to failure at the outset, the system being unable to complete a successful transition to a stable post–contingency state. On the other hand, the system may withstand a disturbance and attain some post–contingency operating state but be unable to respect acceptable steady–state conditions or be subject to slowly developing instabilities such as voltage collapse [Mansour & Kundur 1991].

The social impact of power system shutdowns, whether partial or total, differentiates power systems from any other technology. The National Energy Board of Canada [National Energy Board of Canada 1967], in the Introduction to its *post–mortem* of the 1965 Canada–Northeastern United States blackout, states:

> "The scale of the northeast power failure of November 1965 was immense. The power flow interrupted was some 24 million kilowatts. The initial impact upon the 30 million persons directly affected mounted rapidly as minutes changed to hours without electric service. Concern moderated in most cases when public broadcasts reassured people that the cause was not catastrophic. That such a disastrous event could happen to a large sector of a great and sophisticated industry, and that it could have occurred under much less favourable circumstances, are sobering thoughts indeed."

The ready availability of electric power not only contributes to the pleasure and comfort of our society, it constitutes an essential infrastructure of our civilization [Trotignon et al. 1992]. When central utility service is interrupted, beyond momentary inconvenience or economic loss lies the threat of large–scale social unrest, violence, the breakdown of communications, perhaps loss of life. This can even be the case for service interruptions of relatively short duration, particularly in densily populated neighbourhoods, as was seen in the New York City blackout of 1977 [Corwin & Miles 1978]. The issue of maintaining the integrity and continued operation of a power system, its "survival", is therefore strategic to society as a whole though, for the electric utility, it also makes sense from a strictly business point of view.

The first serious attempts to circumscribe power system security as a discipline in its own right emerged in the wake of the Canada–Northeastern United States blackout of 1965 [Byerly & Kimbark 1974; see Introduction]. This can clearly be seen from the previously cited National Energy Board of Canada document which reads as follows:

> "The immediate result of the blackout was that the power utilities involved ... at once realized it was essential to take a long, hard, fresh look at electric power systems in relation to the security of power supply. ... Are the planning criteria right? Are the design concepts adequate? Are the operations criteria correct? Are operations personnel adequately equipped and trained? Should interconnections between power systems be avoided or increased? The answers to these and to many similar questions are not being found without painstaking studies."

Most authors credit Dy Liacco for laying down the theoretical foundations of power system security in a series of reports and papers published in the late 1960s and 1970s [Dy Liacco 1967, 1968, 1974, 1978]. He defined security in terms of satisfying a set of equality constraints over a subset of the possible disturbances called the "next contingency set" [Dy Liacco 1968]. In fact, as pointed out by Tinguely [Tinguely 1992], several definitions for power system security have been proposed over the years. For example, security has been described alternatively as "an instantaneous time varying measure of the robustness of the

system relative to imminent disturbances" [Fink & Carlsen 1978, Loparo & Fayez 1990, Balu et al. 1992], as "the maintenance of supply (i.e. avoidance of loss of load)" [Stott, Alsac & Monticelli 1987] or as the "ability ... to withstand disturbances ... not result(ing) in large–scale interruptions" [Fouad, Venkataraman & Davis 1991]. The North American Electric Reliability Council (NERC) which provides reliability and security guidelines to all the utilities in North America defines security as "prevention of cascading outages when the bulk power supply is subjected to severe disturbances", as cited by Balu [Balu et al. 1992]. This definition clearly evokes the cascade failure of dozens of utilities within 7 seconds in 1965 as reported by the National Energy Board of Canada [National Energy Board of Canada 1967]. All of these definitions can be brought together as one by regarding power system security as the art and science of ensuring the "survival" of power systems.

1.2 Power systems security in the operations environment: circumscribing the problem

The primary function of the system operator is to provide intelligence at the highest level of a power system's closed–loop control strategy. Operator intelligence closes the control loop, however and whenever necessary, to implement a defensive control strategy, that of overseeing network security and ensuring continuity of service. In a sense, the sole purpose of many on–line monitoring or decision tools is to aid the operator's decision process in order to "establish defensive operating states where no single contingency event (i.e. a single failure) will cause overloads and/or out–of–limit voltages" [Wood & Wollenberg 1984].

From the perspective of the system operator, a network is a dynamically changing entity. Not only does the operator *authorize* planned topological changes, meaning the withdrawal, inspection, maintenance, repair, commissioning and re–energizing of major power system components, he or she reacts, by means of appropriate actions, to assure network security when, and even before, unforeseen events occur in the network. As can well be understood, a network will very rarely be found in its so–called *noble* topology, that is, the original network design as laid down by the system planner. Networks usually find themselves in some reduced topology, comprising fewer elements than the noble network, often termed a *degraded* network. This being the case, how is it

possible to assure power system security when network topology is changing all the time?

Here, we touch upon one of the key functions of operations planners. The basic approach is as follows: for all degraded network topologies likely to be encountered in the power system's operation, the operations planner must provide guidelines to the operator. These guidelines enable the operator to assess system security: In the event of an identifiably insecure state, corrective rescheduling, remedial action or preventive rescheduling are used to adjust the level of security in such a way as to ensure survival of the network in the event of likely contingencies [Stott, Alsac & Monticelli 1987]. The problem is consequently shifted to that of ensuring the security of every possible degraded topology, a problem of combinatorial dimensions [Fouad 1988, p. 1310; Balu et al. 1992 p. 276].

By all accounts, the scale of this problem is immense [Fouad, Venkataraman & Davis 1991; Introduction]. Even if a conservative estimate is made of possible degraded network topologies and of likely contingencies (i.e. taking into account different types and different locations), one begins to see that the point–wise approach, to use a term proposed by [Fishl, Mercede, Wu & Chiang 1988] (i.e. for each identifiable topology, operating point, contingency type and location, one determines whether or not the network is secure), is almost unmanageable. The approach only works because of the considerable domain expertise of a few individuals who have evolved with their networks and are able to downsize the problem, often by means of network–specific heuristics.

Because of the size of the off–line security analysis problem, the industry has adopted another approach: the basic strategy involves moving away from an off–line strategy based on a list of possible degraded networks to the on–line security assessment and control of the real–time network. This reduces the problem considerably, focusing only on the real–time network topology. The modern Energy Management System (EMS) is the power industry's major response to this challenge. As summarized by Stott, Alsac and Monticelli [Stott, Alsac & Monticelli 1987]:

"It performs extensive on–line monitoring, assessment and

optimizing functions for the network, to forestall or correct operational problems while maintaining economy. The economic and social benefits of these functions are intuitively great, although they are not easy to quantify."

This strategy is feasible provided that the response time of security–constrained control software is short enough to be of use on–line to the system operator. Strategies based on contingency analysis and steady–state network analysis (i.e. power flow software), the cornerstone of EMS security assessment and control, are capable of addressing the steady–state aspect of power system security in adequate time frames, usually on the order of 15 minutes, though this is achieved thanks to simplifying assumptions, particularly in the way of contingency definition, selection and evaluation. If one attempts to state the steady–state security analysis problem in its most general form, that of a contingency–constrained optimal power flow (OPF) including some kind of economic dispatch, this appears to be just beyond the grasp of existing technology as "the main computational goals of speed, reliability and flexibility seem to have large areas of mutual incompatibility" [Stott, Alsac & Monticelli 1987]. Though a large consensus exists that theoretical and practical problems have yet to be overcome, others are more optimistic and at least one major effort is under way [Carpentier 1987]. Nevertheless, for a list of likely contingencies, existing strategies successfully identify pre–contingency operating states that ensure acceptable network operation within established post–contingency steady–state voltage limits, transmission line and equipment thermal limits, power–flow feasibility limits and short–circuit limits. The respect of such limits is referred to as steady–state security.

However, as we have seen in the previous section, security assessment must also take into account the power system's capability to realize the transition from pre– to post–contingency steady–state: This cannot be taken for granted. In tightly–knit networks, experience has shown that dynamic and steady–state security limits essentially coincide and a steady–state, on–line approach is often adequate. But for interconnected, loosely–knit, radial, heavily–loaded or longitudinal networks, particularly those including long EHV transmission, dynamic security issues are often more restrictive than steady–state security considerations [Fouad 1988, Balu et al. 1992]. Herein lies the

weak point of the on-line steady-state security approach: There is no guarantee that the network will survive the transition to the post-contingency state. Dynamic security analysis therefore concerns itself with power system security based on the network's transient response [Fouad 1988].

Dynamic security analysis is traditionally an area which requires considerable expertise and the use of complex tools to establish transient or long-term response characteristics. Whether these tools employ direct methods, such as the transient energy function method [Pai 1989, Fouad & Vittal 1991], the extended equal area criterion [Xue, Van Cutsem & Pavella 1988, 1989, 1991] and mode analysis [Chiang, Tong & Miu 1992], or indirect methods such as transient or long-term stability simulation packages [Byerly & Kimbark 1974], they demand orders of magnitude more computing resources than steady-state security tools. For instance, a commercial load flow package may find a network solution within seconds for a typical network, whereas a transient stability simulation can take anywhere between half an hour to half a day, depending on simulation time length and the speed of computation. Dynamic security analysis consequently remains confined to being, at present, an off-line activity though real-time computing capability (i.e. for instance, for execution of transient stability simulations in real time) would render on-line strategies attractive [Balu et al. 1992].

In practice, no utility has the resources to cover all of the possible degraded network topologies at the present time [Akimoto et al. 1989]. To arrive at working guidelines, operations planners rely on heuristic knowledge and experience, usually leading to conservative operating strategies. But what form must these guidelines take? Here, one must not lose sight of a fundamental consideration: The *raison d'être* of power systems is to generate, transmit and deliver electrical power, and power system economics require that electrical facilities be operated at or near their rated capacity. Intuitively, we understand that there is some link between transmission capacity and dynamic security guidelines. Can we show this rigourously? To address this issue, we must first examine the underlying processes of system and operations planning.

1.3 SYSTEM PLANNING, OPERATIONS PLANNING, CRITERIA AND VOLTAGE CONTROL

Traditionally, electric power utilities have dealt with security using deterministic criteria [McGillis et al. 1992]. From a system planning perspective, this means that a network is designed to meet or exceed overvoltage, stability and reliability criteria which, from experience, are held to be so conservative as to cover all but the most improbable contingencies. Criteria are standards that relate to the performance of the system under appropriate steady–state and transient conditions. For example, a six–cycle, three–phase fault with subsequent loss of line after fault clearing is a typical contingency used as a stability criterion; a steady–state temporary voltage capability of 1.5 p.u. is a typical overvoltage criterion; one complete power system shutdown every thirty years is a typical targeted reliability criterion [McGillis et al. 1992].

1.3.1 SYSTEM PLANNING

The system planning process is made difficult by the fact that such criteria are not entirely independent of one another: The choice of a particular mix of criteria results from a judicious balance between competing economic, environmental, geographical and technological constraints [Galiana, McGillis & Marin 1992]. This signifies that system planning is inherently an iterative process where sensitivity analyses play a crucial role. Stated in the simplest terms, the system planner seeks to meet some secure transmission capacity subject to established criteria, and network design is the variable. This is a far–reaching statement: It implies that transmission capacity is topology– and criteria–dependent. Of course, the process must also yield a network design which is "demonstrably operable" [Galiana, McGillis & Marin 1992], a term usually equated with design simplicity.

1.3.2 OPERATIONS PLANNING

For system security to be assured in the operating environment, it follows that operations must be consistent with the original network design philosophy and criteria. In other words, the original design criteria must apply to every degraded topology of the noble network. But the operations planning problem differs from that of system planning in that the degraded topologies and criteria (i.e. such as overvoltage) are known, and transmission capacity therefore constitutes the only remaining degree of freedom. Consequently, the guidelines provided to the system operator are transmission capacities for the progressively

degraded topologies, or as power engineers would say, for increasingly *weakened* networks. When referring to transmission capacity limits obtained in relation to *post-contingency* thermal, load-flow feasibility or steady-state voltage criteria, these are called steady-state security limits. When referring to limits obtained in relation to transient or long-term stability criteria, the resulting limits are termed dynamic security limits.

Of course, the transmission capacities provided to the system operator are values which will guarantee that a given degraded topology is secure for every one of a list of likely contingencies, including the worst. As we have seen in section 1.2, individual power transfer limits must in principle be found for each likely contingency and for each degraded topology, and the most restrictive power transfer limit, whether resulting from steady-state or dynamic security considerations, is the security limit. Generally, the security limit is the most restrictive power transfer limit which identifies, in turn, the "worst contingency". The resulting security limits are usually compiled in the form topology-dependent tables [Dénommé & Vincelette 1984, Fouad, Venkataraman & Davis 1991] though decision trees have also recently been proposed [Wehenkel, et al. 1989, 1993].

1.3.3 The concept of security margin

A frequently encountered concept in dynamic security analysis is that of security *margin*. This relates to the difference between *allowable* and *maximum* transmission capacity: The operations planner always provides guidelines (i.e. *allowable* transmission capacities) that are less than the *maximum* obtained by simulation in order to provide for uncertainty in the data, the topology and in the operating environment (i.e. load distribution, load models, climactic, etc.), and thereby ensure the robustness of the network. For example, at Hydro-Québec, such margins are usually on the order of 3% of the limit. The *operating* margin relates to the distance between the actual real-time power transfer and the allowable transmission capacity, as defined by the security limit. We will return to the concept of margin in Chapter 5.

1.3.4 Voltage control

In section 1.2, we saw that voltage limits are explicitly considered in the statement of the steady-state security problem. From an operating perspective, due to the difficulties related to performing sensitivity studies in dynamic

security analysis (i.e. lengthy simulations must be executed for each new value of a given parameter), maintaining an appropriate voltage profile is considered a prerequisite of power tranfer: Under normal steady–state conditions, security limits are considered firm provided that system voltage is held within certain pre–determined bounds. This is a reasonable assumption in that reactive–compensation elements with adequate redundancy and appropriate controls usually ensure that this is indeed the case, particularly on EHV systems. System planning criteria offer guidelines for the determination of adequate reactive power resources to deal with such issues as steady–state generation and load, voltage profile, the maximum rate of system load increase or decrease, various import or export conditions or system restoration [Hydro–Québec 1990].

In fact, recent research tends to show that reactive power management for voltage control may be the essential problem in system planning. McGillis et al. [McGillis et al. 1992] have obtained results which tend to confirm that many conflicting constraints, such as stability and overvoltage, can be reconciled by a judicious mixture of series, shunt and dynamic compensation, particularly for long transmission systems. For example, to maintain a flat voltage profile (FVP) at all times, significant amounts of static VAR compensation are required. Likewise, the control of temporary overvoltage (TOV) on very long lines requires a considerable amount of series compensation. As summarized by McGillis, "... when these two requirements are fulfilled, namely FVP and TOV, there is sufficient reactive power in the system to cope with even a three–phase fault. In a sense, the stability performance comes as a fringe benefit: It is an output of good design rather than an input".

Regardless of a utility's voltage control philosophy, the operator's function is strategic, even though a good philosophy will attempt to minimize the number of switching operations to avoid circuit breaker deterioration. For example, if the system voltage profile is unsatisfactory or drifts beyond established operating limits, the strength of the system is affected, and the attendant security level is compromised [Stott, Alsac & Monticelli 1987]. In such an event, the operator must either initiate some appropriate corrective action (i.e. switching of lines, shunt inductances or capacitors; redispatching of generation, etc.), or restrict the maximum permissible power transfer (i.e. reduce the security limit). Such

action ensures that voltage returns within acceptable emergency limits after the "worst contingency".

1.4 Dynamic Security Limits and Criteria: Normal versus Extreme Contingencies

From an operating perspective, power–transfer–type dynamic security limits are convenient security indices in that power flows are readily measured at all points on the network. We have seen that dynamic security limits are topology– and criteria–dependent: one might say that dynamic security limits translate criteria into meaningful and, more importantly, measurable quantities. One might even go a little further and state that the basic problem in dynamic security revolves around the issue of equating abstract criteria to appropriate network quantities: The scale of the dynamic security problem follows directly from the abstractness of the criteria and the inappropriateness of our tools due to a lack of fundamental knowledge.

In the light of such considerations, it is essential that proper care be exercised at the outset in choosing appropriate dynamic security criteria as these can have far–reaching effects. Consider the following: experience has shown that, generally speaking, dynamic security limits determined on the basis of a single–line–to–ground fault will be higher than those obtained from a three–phase fault of equal duration. However, there is a hidden price to these higher limits: the former criterion may result in a network reliability index of, say, 3 system shutdowns in ten years whereas the latter may provide a more reliable network by a factor of ten [McGillis et al. 1992]. In other words, increased network robustness is a direct consequence of establishing transmission capacity on the basis of more stringent criteria: The network is "capable of coping with a much wider range of adverse events" [McGillis et al. 1992] and this in turn reflects on the power system's reliability. Expressed in different terms, the short–term gain of operating at higher transmission capacities can be offset by a long–term drop in reliability. This clearly illustrates that network security is relative to a given set of criteria and that high level criteria are often interdependent.

The dynamic security criteria used in system and operations planning generally conform to established industry standards, primarily to ensure the reliability of interconnected systems and protect them from cascade failures. For

example, the North American Reliability Council (NERC) provides criteria guidelines to all utilities in North America through its four constituent Councils [McGillis et al. 1992]. One of these, the Northeast Power Coordinating Council (NPCC) provides criteria for all electric power utilities in New England, Ontario and Québec. These organizations provide criteria to cover normal contingencies, extreme contingencies and system restoration [Hydro–Québec 1990].

1.4.1 NORMAL CONTINGENCIES

Power system security limits constitute the first line of defense in a power system's overall security strategy. Dynamic security limits are defined with respect to normal contingencies which correspond to probable scenarios of unexpected events occurring on the network. A normal contingency is usually defined as the loss of any element in a power system (i.e. line, transformer, generating unit, etc.), either spontaneously or preceded by a single– or three–phase fault [Hydro–Québec 1990]. This is often referred to as the *N–1* criterion. A power system must be capable of withstanding a normal contingency, occurring at any location in the network, without interruption of service to any customers and remain within emergency operating limits.

Because of inadequate knowledge regarding the relative severity of different types of contingencies in different networks, the criteria adopted by a utility can include a list of normal contingencies. It is instructive to review what one typically considers as normal contingencies in one specific utility [Hydro–Québec 1990]:

> "a) A permanent three–phase fault on a generator, transmission circuit, transformer or bus section with normal fault clearing.
> b) Simultaneous permanent phase to ground faults on different phases of: Each of two adjacent transmission circuits on a multiple circuit tower, with normal fault clearing. If multiple circuit towers are used only for station entrance and exit purposes, and if they do not exceed five towers at each station, then this condition is an acceptable risk and therefore can be excluded. Other similar situations can be excluded

on the basis of acceptable risk, provided that the System Design Coordinating Committee and Operating Procedure Coordinating Committee specifically accept each request for exclusion.

c) A permanent phase to ground fault on any transmission circuit, transformer, or bus section with delayed fault clearing.
d) Loss of any element without a fault.
e) A permanent phase to ground fault on a circuit breaker with normal fault clearing. (Normal fault clearing time for this condition may not always be high speed.)
f) Simultaneous permanent loss of both poles of a direct current bipolar facility.
g) The failure of a circuit breaker associated with a special protection system (SPS) to operate when required following: loss of any element without a fault; or a permanent phase to ground fault, with normal fault clearing, on any transmission circuit, transformer or bus section.
h) With all facilities in service, (i) the simultaneous loss of two parallel transmission lines following a permanent single phase fault on one of the lines, with normal fault clearing, or (ii) a permanent single phase fault on a transmission line with delayed fault clearing during the time that one circuit breaker is unavailable."

As present knowledge does not yet permit us to gauge the relative severity of, say, contingency a) with respect to contingencies c) or f), such lists are indispensable. However, in most situations, long–time network–specific experience reduces the search–space for the most severe normal contingency type and location.

1.4.2 Extreme contingencies

System and operations planners also study the effects of extreme contingencies, such as multiple incidents or cascade tripping, in order to establish operating strategies which will reduce their frequency of occurrence and, most importantly, to limit their consequences: The primary objective of such studies is to provide means of avoiding a complete system shutdown. Special Protection Sys-

tems (SPS) such as generation rejection, shunt reactor switching, load–shedding systems or other such automatic controls constitute the main arsenal of measures used to mitigate the consequences of extreme contingencies. Generally speaking, good design practice avoids the use of SPS to maintain system integrity during normal contingencies. The use of SPS is only considered acceptable as a last resort to guard against the complete loss of the system in the event of extreme contingencies.

To further distinguish extreme contingencies from normal contingencies, it is once again instructive to consider the following list of extreme contingencies in the aforementioned utility's criteria [Hydro–Québec 1990]:

"a) Loss a generating station or part of a generating station.
b) Loss of several lines emanating from a substation, including the loss of all of these lines.
c) Loss of all lines in a corridor.
d) Multiphase short circuits with delayed clearing.
e) Loss of a major load centre.
f) Simultaneous loss of all DC interconnections caused by a disturbance outside the system.
g) Failure of an SPS whose operation is needed following the occurrence of one of the normal contingencies.
h) Complete or partial operation of an SPS during a contingency or under circumstances for which it was not intended to operate.
i) Loss of series capacitor banks on all parallel lines connected to the same substation, after a fault resulting in the loss of one or several of these lines."

1.4.3 System Restoration

In the event that the system has been lost, restoration strategies attempt to provide safe, high–speed load pickup while providing adequate facilities to effect all the necessary switching: The system must be restored without jeopardizing system or component security, whether from a steady–state or dynamic security perspective.

1.4.4 Criteria and the Utility: Strategic Choices

Notwithstanding the existence of guidelines for interconnected systems provided by such bodies as NERC and NPCC, the issue of criteria is strategic in that a utility accepts to employ certain criteria by choice, not because they are imposed in any way. A corollary to this is that the choice of criteria has long–term consequences on the reliability and economics of a power system, though optimal choices ultimately depend on the overall business environment of each specific utility. For example, the choice of an $N-1$ criterion employing a single–line–to–ground (SLG) fault will appear attractive from the perspective of short–term economic considerations though such a choice will result in a lower long–term reliability index than if the more severe three–phase fault (i.e. of equal duration) were used. Additionally, if a utility having adopted an SLG fault criterion is interconnected with other utilities which employ the three–phase fault $N-1$ criterion, the latter will prefer to limit interconnection capacity on the basis of their own criteria, even though the interconnection equipment rating may be higher. Generally, if a utility wishes to upgrade its criterion to a three–phase fault, this translates into an overall reduction of system capacity unless massive reinforcement of the transmission system is undertaken in the form of series compensation, static VAR compensators (SVCs), synchronous condensers (SCs), higher voltage levels, etc. [McGillis et al. 1992].

1.5 Transfer Limits, Worst Contingencies and Dynamic Security Limits: Definition of Terms

Having defined a number of fundamental concepts related to dynamic security, we are now in a position to provide a rigourous definition of such critical terms as "transient stability transfer limit", "long–term stability transfer limit", "dynamic security limit" and "worst contingency" which are used throughout this thesis.

A transient or long–term stability transfer limit is the pre–contingency power transfer value, associated with a transmission line section or corridor, for which the application of a normal contingency at some location within said line section or corridor will result in the highest power transfer consistent with the fulfillment of some predetermined acceptance criteria. In other words, this represents the maximum power transfer capacity of the network, or of some part of the network, consistent with its capacity to sustain a specific normal contin-

gency. A typical acceptance criterion might state that the network must be stable and that post-contingency voltage must not be less than .95 p.u. after 8 seconds at all load buses.

The distinction made between transient and long-term stability simulations resides essentially in the length of the simulation, though there is also some difference in the focus regarding the phenomena to be modeled and observed. As defined in [IEEE Task Force on Terms & Definitions 1982], short-term or transient stability is limited to study of the behaviour of a system for several seconds following a large-scale disturbance (i.e. a fault, change of topology, etc.); long-term stability relates to the study of system behaviour beyond several seconds, usually within a time-frame appropriate for the modelling of load cycles, tranformer tap behaviour and shunt reactor switching. Long-term stability is currently an important tool for the determination of voltage-collapse-dependent transfer limits [Gao, Morison & Kundur 1992].

A transient or long-term stability transfer limit is always given in association with the contingency location. The location of the normal contingency which gives the lowest power transfer limit is defined as the worst contingency. The power transfer limit associated with the worst contingency is called the dynamic security limit.

1.6 Dynamic Security Assessment and Control: A Literature Review

Dynamic security assessment and control are clearly very important to power system operation. However, in section 1.2, we have seen that the scale of the problem is such that existing off-line strategies leave considerable room for improvement. Whether viewed from 1) the lengthy, expert-based processes that have evolved within the industry over the last few decades to build dynamic security limit tables, 2) the tools used to perform specific tasks within these processes such as power flow, short-circuit, energy methods or transient and long-term stability software, or 3) the fundamental domain knowledge that guides experts within the processes, there is no need to stress that dynamic security analysis provides a wealth of extremely attractive and timely research opportunities.

Because steady-state security analysis strategies and technologies have

gravitated relatively quickly to the on–line EMS environment, they have exercised considerable influence on their counterparts in dynamic security analysis. It is therefore appropriate to begin this literature review by examining existing practices in steady–state security.

1.6.1 Steady–state security strategy: The point–wise approach

Since Dy Liacco's well–known papers describing security in terms of normal, emergency and restorative states [Dy Liacco 1967, 1968], considerable effort has gone into the study of steady–state security, whether from a strategic perspective [for example Cihlar et al. 1969, Galiana, Swcheppe & Glavitsch 1974], from an off– or on–line *process* point or view [Balu et al. 1992], or in consideration of the tools required to support various high–level approaches [as reviewed by Stott, Alsac & Monticelli 1987].

Existing practices and trends in dynamic security analysis find many of their roots in the largely successful industry–wide steady–state security developments, culminating in on–line security evaluation. At the heart of this approach is the operations control center, more commonly known as the energy management system (EMS), which is seen to have three functions [Wood & Wollenberg 1984]:

1. System monitoring.
2. Contingency analysis.
3. Corrective action analysis.

System monitoring is generally considered the most important of the three, probably due to the view that quality information in the hands of the system operator will go a long way to improving the highest–level security control loop. For instance, thanks to a better appreciation of network voltage conditions, the system operator can provide more effective voltage control. Among others, this has motivated the development of 1) telemetry systems, for measurement and data transmission, 2) data bases for gathering the data, 3) software for data processing and display, and alarm processing, and 4) much work on state estimation, where the security of the measurement system itself is an important issue [see for example Crainic, Horisberger, Do & Mukhedkar 1990].

Nevertheless, the published literature is richer in the area of *contingency*

analysis, which includes definition, selection (i.e. screening and ranking) and evaluation of contingencies. This is due to the fact that contingency analysis is extremely time–consuming, with a significant and very unwelcome impact on computer sizing: Greater computer hardware size and power (i.e. memory and performance) translates into greater cost, and smarter contingency analysis can reduce this cost.

The basic idea in contingency analysis is simple: Each credible contingency must be simulated on the network at hand. In section 1.2, we have seen that this is referred to as the *point–wise* strategy [Fischl, Mercede, Wu & Chiang 1987] of deliminating the security region, due to its emphasis on individual trials (i.e. one point at a time) in the space of topologies, network operating points and contingencies. Because time constraints are critical in the on–line environment, one must somehow prune the search space of contingencies for the worst contingency. Typically, this is done, first, by defining a list of credible "worst" normal contingency types and locations: This list varies with topology and load, and may include secondary (i.e. state–dependent) switching, resulting in further contingencies. Then, contingency selection reduces the size of the list by entering a *screening* process, which attempts to identify and eliminate cases having no violations. This is usually accomplished by means of approximate power system models and fast, limited–accuracy, computational techniques [in particular, those proposed by Galiana 1984; Brandwajn 1988, 1989; Lauby 1983, 1988]. This is followed by a *ranking* of cases on the basis of severity criteria [see for example Ejebe & Wollenberg 1979; Irissari & Levner 1979; Mikolinnas & Wollenberg 1981]. Ranking strategies may also take into account "masking" effects due to identical rating of cases which are not, in fact, equally severe [Halpin, Fischl & Fink 1984]. Finally, the cases are evaluated using a full power flow in decreasing order of severity, yielding the desired security limits [Stott, Alsac & Monticelli 1987].

The objective of *corrective action analysis* is to provide the system operator with alternative corrective or preventive control strategies, respectively for violations in the immediate real–time network or for the worst–case anticipated post–contingency network. For example, a simple type of corrective action involves redispatching generation among a set of power stations, causing power flows to change, thereby altering line loadings [Wood & Wollenberg

1984].

The present trend is towards greater integration of these tasks, notably by means of steady–state security–constrained OPF [Carpentier et al. 1989]. This, coupled to efforts in adapting vector processing [Anderson & Wollenberg 1992] or concurrent processing [Crow, Tylavsky & Bose 1991], constitute the leading edge of research in this area. However, some have also begun to focus their efforts on expert systems [Liu & Dillon 1990, Tinguely 1992].

In the light of the success of the point–wise approach in steady–state security analysis, developments in dynamic security analysis have tended to follow a similar pattern, though considerable economic and technological obstacles remain before on–line power transfer limit determination, not to mention transient or long–term stability–dependent contingency analysis (for security limit determination), becomes tractable. We shall see this in 1.6.2 as we survey existing industry practices, also based on the point–wise approach and indirect methods. In 1.6.3, we will describe the traditional alternatives (i.e. the direct methods) and show why they have not as yet provided a serious challenge to replace present practices. However, as we shall see in 1.6.4, direct methods may yet emerge as serious candidates in contingency filtering strategies. We complete this literature survey in 1.6.5 by showing how artificial intelligence and expert systems have contributed to addressing important aspects of dynamic security analysis processes, and to position emerging developments in such areas as pattern recognition and neural networks.

1.6.2 Present industry practice: The reliance on indirect methods

Notwithstanding the considerable efforts deployed since the middle 1960s, power system stability continues to elude attempts to characterize its fundamental aspects in a compact, definitive way. In describing recent research performed in dynamic security, Liu [Liu 1990] remarks that "The amount of literature on power system security is large: most of the work is concerned with the development of numerical algorithms".

Though this observation simplifies to a certain extent, one does indeed observe that the mainstream of research in dynamic security analysis abounds with papers treating highly specific aspects of the point–wise approach, often

bringing depth to specialized areas, but offering little horizontal perspective. For example, such areas as component modeling [Sauer & Pai 1991], transient stability time–simulation software and stability enhancement through smarter and faster controls [Byerly & Kimbark 1974, Langevin 1986, Wang et al. 1992] have been the focus of considerable research effort. Some have attempted to characterize stability from some unique or practical perspective, for instance, by estimating system damping from statistical analysis techniques [Torre & Bushner 1987] or by studying the Fourier transform behaviour of tie–line flows [Saitoh, Toyoda & Kobayashi 1991], yielding very interesting results. Advanced control and protection systems [Phadke & Thorp 1991, Shoults & Jativa 1991, Venkata, Damborg & Jampala 1991], better protective relaying [Horowitz 1980, more recently Thorp & Phadke 1991], and the use of parallel or concurrent computing [Crow, Tylavsky & Bose 1991, Chai & Bose 1992] have also emerged over the years. In fact, most of the references cited here (i.e. notably those from 1991) review the state of the art in their respective areas and collectively refer, in turn, to hundreds of papers. But how does the industry integrate these results to ensure the real–time dynamic security of power systems?

1.6.2.1 THE BASIC APPROACH

Fouad [Fouad 1988], in an IEEE working group report on industry–wide dynamic security assessment practices, describes the current approach as follows:

> "... off–line studies are performed for different initial operating conditions, for a prescribed sequence of events or contingencies. From these studies, "safe" operating levels are arrived at for a variety of system conditions. These are often given in terms of a critical system operating parameter such as: the loading of a certain power plant, the power flow at critical transmission interface, the voltage at a given bus, and so on."

In association with others, the author has previously described existing operations planning practices at Hydro–Québec [Marceau, Mailhot & Galiana 1992]. Consider Fig. 1–1. Given a starting network topology, the base network Nb, the operations planner first selects a set of probable degraded networks, Nd_i. Each Nd_i consists of the base network after the loss of an EHV line section (i.e. a 735 kV line on the Hydro–Québec grid), a complete EHV line (several

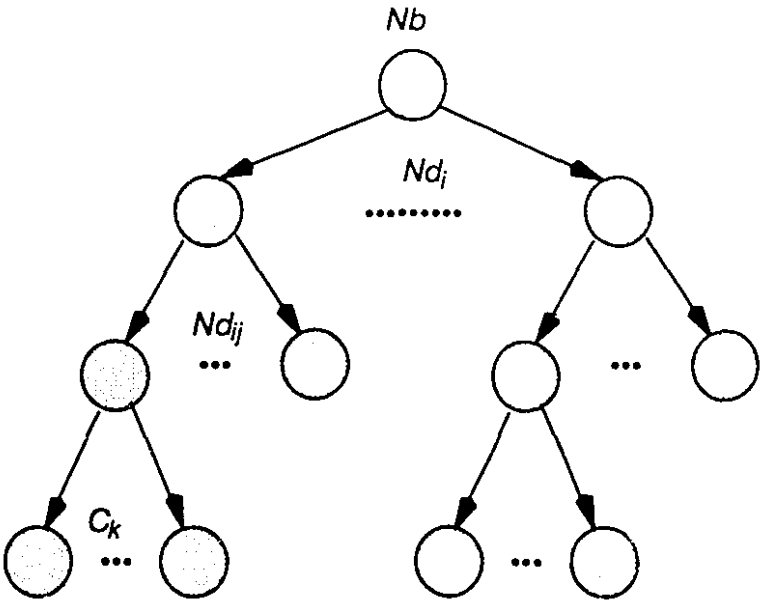

Fig. 1–1. The logic of dynamic security analysis in operations planning.

serially–connected line sections between a major generation centre and the load), or a combination of line sections or lines in adjacent transmission corridors. For each Nd_i network, the planner then assumes unavailability of combinations of voltage support components (SVCs, synchronous condensers or shunt reactors) or of tie–lines between corridors; this gives the Nd_{ij} networks. The idea is to anticipate all possible operating conditions arising from scheduled maintenance or unexpected equipment failure. The Nd_{ij} networks are those simulated for each of a set of normal contingencies C_k: Transient and long–term transfer limits are obtained for each combination of Nd_{ij} and C_k, and dynamic security limits are, in turn, derived from these results.

The studies which follow are therefore parameterized in terms of the unavailability of some subset of facilities. For each case, safe conditions are tabulated in separate tables and, during on–line operation, a given network configuration is used as a key to find and access the most applicable table [Avramovic & Fink 1991]. The on–line use of these tables consists of finding a previously studied network that is closest in terms of configuration and loading, and secure operation of the network is ensured provided that system operators maintain power transfers within these limits.

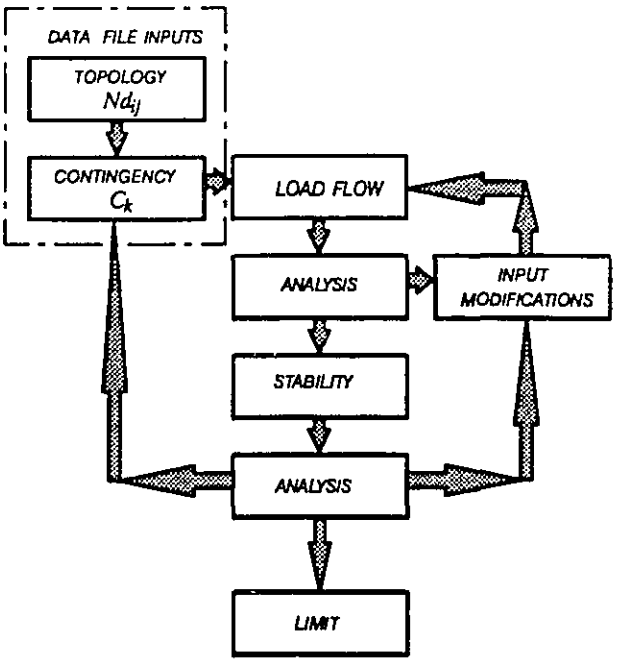

Fig. 1–2. The power transfer limit determination process.

1.6.2.2 STABILITY TRANSFER LIMIT DETERMINATION

As can well be seen, transient and long–term stability transfer limit determination is the key to this approach. However, finding such limits is a complex, iterative process which requires the execution of many time–domain stability simulations and considerable expertise.

Consider Fig. 1–2. To perform a single power transfer limit search, one must first execute a load–flow software package for a given topology, analyze the results, carry out manual modifications to the data input and repeat the same procedure until a satisfactory steady–state case is found. This initial step is not trivial: even such deceptively simple areas as load flow analysis and correction can be so complex as to warrant optimal load flow or expert system tools on their own [Cicoria, Migliardi & Marannino 1993]. A satisfactory steady–state case is then used to initialize the network for a transient–stability simulation, also manually initiated by the user. When the simulation is finished, one will often resort to yet another package to extract results, perform transient or long–term stability analysis, apply acceptance criteria, determine what the next step will be (i.e. increase or decrease power transfer in the faulted corridor), modify load–flow software inputs accordingly and re–enter the process. This is repeated un-

til the required limit is found to the desired accuracy. In principle, to find the security limit associated with a transmission line section or corridor, one must repeat this process for different contingency types and locations (i.e. on the line section or corridor, see 2.4.1) until the most constraining (i.e. lowest) transfer limit has been identified.

1.6.2.3 DYNAMIC SECURITY ANALYSIS AS A CLASS OF PROBLEMS OR PROCESSES

Further study of Fig. 1–2 shows that dynamic security limit determination is a complex *process* which consists of many smaller, individual problems. These problems include finding 1) an appropriate power flow for a given topology from which to begin a transfer limit determination process, 2) the security status of the topology (i.e. is it stable or not?) for a given set of injections and loads with respect to some contingency, 3) the power transfer limit of some part of the network with respect to a contingency applied at some location, 4) the security limit of some part of the network with respect to some contingency and 5) the associated worst contingency location. Often, they must be addressed within a larger framework, for instance, that of optimizing security limits as a function of some network parameter (i.e. commonly referred to as sensitivity studies). Viewed from the perspective of existing practices, dynamic security analysis can be seen as a class of problems where high–level strategies and tools have yet to be developed to deal efficiently with each particular problem type.

To fully appreciate the impact of this statement, one need only consider that such a large number of combinations of Nd_i and Nd_{ij} networks, in addition to normal contingencies C_k, must ideally be considered that typical operations planning departments are unable to consider them all [Akimoto et al. 1989]. There are two main reasons for this.

1) Dynamic security analysis processes are extremely time–consuming, if only from the point of view of expert analysis;
2) The analysis of every conceivable degraded topology is a problem of combinatorial dimensions.

Let us consider the first point. Though the computation times of load flow, transient and long–term stability simulations have dramatically decreased in recent years, the *process* remains highly time–intensive due to expertise–re-

lated tasks. Data collection, input validation, output post–processing and result analysis still require hours whereas a single 15–second stability simulation of a 700–bus, 1000–line network takes less than an hour on a SUN Sparcstation 2. Human error, incorrect data, insufficient expertise and the boredom of routine tasks all contribute to this problem. Notwithstanding the need for extensive computation, one is forced to conclude that the process loop is rendered inefficient because of the human element which closes the loop.

Considering now on the second point, because of the number of topologies which can conceivably be examined, expertise is used to scale the problem down so that the dynamic security limit of only a small set of "umbrella" topologies is explicitly found. Extensive analysis of this reduced set, focusing primarily on various combinations of voltage support equipment for the highest voltage network, forms a basis for the heuristic estimation of dynamic security limits covering all Nd_{ij} networks. But there is a price to pay: small sample sets translate into lack of precision, resulting in more conservative security limits and higher operating costs.

Apart from this, other issues must be resolved before the system operator can make good use of these results. The sheer number of tables generated by the off–line studies makes a manual search an arduous task for an EMS operator. Some utilities have developed semiautomated procedures that alleviate this burden, such as entering the tables into a data base that is easily accessed by the EMS: Once an appropriate search logic is in place, the EMS is able to continuously verify the tranfer limits. However, the packaging of off–line studies into appropriate tables and the design and implementation of table access logic are still performed manually [Avramovic & Fink 1991].

Crevier [Crevier 1978] was the first to address the key issue of accelerating dynamic security analysis processes. Using a simplified transient stability software, a simulation–based search algorithm was proposed to determine points on a surface in the space of injections and loads defining the security region. It was hypothesized that the surface was either a hyperplane or a quadratic surface, and a method was proposed to determine the coefficients of each type of surface. The two types of surfaces were compared with the actual security region obtained by simulation on a three–generator system, and good agreement was

observed. The method was also used to optimize power system stabilizer settings, again with good results [Nourmoussavi & Crevier 1979]. Unfortunately, this work did not enjoy wide acceptance at the time and was not pursued.

Clearly, before dynamic security analysis can gravitate to the on-line environment, as in the case of steady-state security evaluation, the processes of dynamic security analysis must be mechanized in some way. A prerequisite to this however, independently of hardware or software considerations, is that one must formulate dynamic security analysis in terms of its constituent goals and processes, so that these can in turn be mechanized within the context of an effective, coherent on-line strategy.

1.6.3 DIRECT METHODS OF STABILITY ANALYSIS

Since Athay, Podmore and Virmani proposed a transient energy function in the late 1970s [Athay, Podmore & Virmani 1979], important advances have taken place [Fouad & Vittal 1988, Pai 1989] and this approach has almost become synonymous with direct methods. However, in parallel with this work, a clever extension of the classic equal area criterion has been proposed for transient stability analysis of multimachine systems [Xue, Van Cutsem & Pavella 1988, 1989, 1991] which has aroused interest.

The energy function method describes the disturbance-induced change in energy of the power system as an analytical expression in terms of power system variables and parameters. Analysis of this expression for different topologies and contingencies, based on Lyapunov stability criteria, permits the explicit identification of a security region. Stability thresholds (i.e. limits) can be determined from such criteria and performance indices identified to gauge the relative severity of particular cases [Fishl, Mercede, Wu & Chiang 1988]: If the value of the transient energy function is less than the threshold, the system is stable [Fouad & Vittal 1988]. From this approach, security regions can be viewed as hypervolumes in the space of injections (i.e. generation and bus loads) [Wu & Tsai 1983]. Also, because energy functions describe a continuum, these methods also permit the explicit determination of security region sensitivity to changes in the state variables.

Though energy function methods are faster than the indirect methods, they

have been slow to gain acceptance for a number of reasons. First, nonlinear or active control elements and high–speed control strategies (i.e. SVCs, static excitation systems and complex control systems) have been difficult to treat, though this appears to be changing [Hiskens & Hill, 1991]. Second, one must simulate fault–on power system dynamics so that the energy can be computed upon clearing the fault: The most expedient way of doing this is by means of a transient stability software package. Third, an important hurdle remains the determination of the critical energy (i.e. the threshold) at which the system is unstable. Rigourous approaches have traditionally generated conservative values for critical energy though heuristic reasoning has assisted in improving this [Fouad & Vittal 1988]. Fundamental work is also emerging, particularly in the area of mode analysis, which promises to improve critical energy estimation [Chiang, Tong & Miu 1992].

Finally, because of the preceding points, the method not been frequently applied to existing networks and there is a lack of data comparing their performance to state–of–the–art time–domain simulation packages [Vittal 1991]: This is an important point in a highly conservative industry, though this is also changing [El–Kady et al. 1986].

The extended equal area criterion is an alternative direct method, based on the classical equal area method originally developed for the one–machine–infinite–bus problem [see Introdution, Byerly & Kimbark 1974] but extended to deal with large networks. Using heuristic knowledge and clever insight, a multimachine power system is replaced by a two–machine dynamic equivalent further amenable to the one–machine–infinite–bus problem. The stability problem is then reduced to an algebraic equation, obtained from the equal area criterion, and this equation is used to derive "one–shot" stability analysis strategies [Xue, Van Cutsem & Pavella 1988]. Here again, we have a very fast tool, and first and second swing stability results appear to be excellent. However, as in the previous method, because of the difficulty of treating nonlinear or active, high–speed control elements, the method yields pessimistic results for power systems incorporating such devices and acceptance has consequently been slow. In addition, some experts feel uncomfortable with an approach which, due to its heuristic assumptions, provides a one–way transformation of the system for the overall stability picture: Though this is in itself highly valuable, the method is

as yet unable to provide a return mapping of stability information onto individual machines [Xue, Van Cutsem & Pavella 1988; see discussion by Venkata].

Modal or eigenvalue analysis is yet another method of obtaining insight into stability behaviour, and its main strength is that it can explicitly identify unstable power system modes [Elgerd 1967, Desoer 1970, Elgerd 1971]. The fine paper by Gao [Gao, Morison & Kundur 1992], though not having anything to do with transient stability, serves as a powerful reminder that a return to first principles can sometimes be extremely fruitful: In this case, modal analysis is applied to the steady–state (or long–term) voltage stability problem. Unfortunately, power system stability modal analysis has traditionally been confined to small–signal linear approximations of a power system [Elgerd 1971]. To a certain degree, this is a disappointment in that current industry practices favouring the point–wise approach would clearly benefit from system–wise approaches. Nevertheless, attempts have been made to apply modal analysis for the optimal control of power systems, most notably in the design of an adaptive voltage regulator [Langevin 1987]. The regulator settings are computed on the basis of an optimization strategy which consists of minimizing the energy expended by the system to reach equilibrium. The approach is very promising but requires further development.

1.6.4 Direct versus indirect methods

As a result of the shortcomings of existing direct methods, indirect methods, in the form of time–domain simulations employing any one of several commercial grade transient/long–term simulation packages, are the industry standard for determining dynamic security limits. They are capable of treating large–scale disturbances, complex non–linear and active network elements and elaborate contingency scenarios, including protections. Their widespread acceptance is clearly seen in the report given by the IEEE Working Group on Dynamic Security Assessment Practices in North America [Fouad 1988]. As a consequence of this, the point–wise approach is almost universally used for determining dynamic security limits though human resource–based limitations impose severe constraints.

However, the point–wise approach is a suitable strategy for the on–line environment, as in the case of steady–state analysis, provided that secure trans-

mission levels or operating margins can be identified within an acceptable time–frame. Needless to say, computation speed is an issue, particularly in relation to time–domain simulations, but the entire family of dynamic security limit search *processes* must also be *mechanized and optimized* before the approach can be considered viable. With this in mind, direct methods of stability analysis, previously viewed as competitors to the more traditional indirect (i.e. time simulation) methods, are now being viewed – more pragmatically – as candidates to filtering strategies in limit–search processes [Fouad & Vittal 1991]. In fact, as we shall see in the next section, some expert systems have proposed a division of tasks along these lines, using the extended equal area criterion or transient energy function as filters within different contingency analysis strategies [respectively Xue 1988; Fouad, Venkataraman & Davis 1991].

1.6.5 Expert systems and artificial intelligence

In order to improve the performance of the off–line dynamic security analysis process or to pave the way for on–line strategies, a number of approaches have attempted to mechanize a variety of human activities in dynamic security analysis by means of expert systems, artificial intelligence algorithms, or some mixture of algorithmic and knowledge–based (i.e. hybrid) systems. These efforts have been directed at both low– and high–level tasks within the limit–search process, and have only recently begun to emerge [Xue, Gao, Zhu & Liu 1989].

Wehenkel [Wehenkel, Van Cutsem & Pavella 1989] has proposed a unique approach which is essentially as follows: Large numbers of simulations are performed off–line, covering many contingencies and degraded topologies; the stability status of each is then noted, in addition to various user–defined stability attributes such as bus voltage, generation, etc. (the "learning set"). Applying the method of inductive learning (i.e. learning by example), an "artificial intelligence algorithm" then attempts to build a decision tree for each topology based on the selected attributes which, in turn, can be used in the on–line environment. The decision tree is then integrated within an energy management system (EMS) in order to monitor the stability of the real–time topology. An example of this is provided for a 14–bus topology of the Greek power system which gives plausible results.

The advantage of this approach is that the algorithm generates decision trees independently of human intervention. However, if the learning set of off-line simulations is too small, the algorithm is sometimes incapable of determining the stability status of all operating states generated by the decision tree: A number of states is thereby classified as unknown, though uncertainty diminishes with increasing numbers of simulations. These numbers can be very large: in one example, 25,000 load flow runs and 9,000 critical clearing time computations (using step-by-step simulations) are reported [Wehenkel et al. 1993]. Notwithstanding these weaknesses, the method has been tested on part of the Électricité de France (ÉDF) system with apparently encouraging results [Wehenkel et al. 1993].

Xue [Xue 1988] has proposed a framework for a hybrid expert system including transient stability assessment and transient stability preventive control. This system is reported to integrate existing algorithm-oriented methods with a hierarchical knowledge base. In particular, the fast extended equal area criterion acts as a screening tool, and step-by-step transient stability simulation deals with sophisticated models when these are required. Subsystems on the low level (i.e. using modular, dedicated expert systems) aim at different domain objectives, the high level with metaknowledge performs a coordination role and uncertainty is accounted for using confidence factors. A number of languages were considered (i.e. C, LISP and FORTRAN) and FORTRAN was used in consideration of speed issues (i.e. for simulation software and expert system construction). Of course, this solves the major hurdle of integrating the different software modules, while leaving open the larger issue of using existing, tested commercial simulation software. In 1989, the system was described as being in the process of being improved [Xue, Gao, Zhu & Liu 1989] to include decision support capability in system and operations planning. The conceptual design of this system is extremely interesting but no account has yet been published of the system's performance.

Akimoto [Akimoto et al. 1989] describes a "transient stability expert system" for the system planning environment but which may also be useful in operations. The transient stability problem is represented as an "iterative design process with some similarity to closed-loop feedback control" though "review and action by an engineer in the stability analysis process is still required

at this stage". The paper describes an expert system study organization environment where a particular goal can be defined, certain system conditions and contingencies identified, data validated and criteria established. Following this, depending on the goal, the system generates a list of worst contingencies, screens and ranks them using the transient energy function method, and then initiates time simulations on this basis. It appears that stability analysis is performed interactively with an expert user. The system also performs problem cause analysis in an interactive mode with the user, primarily on the basis of transient energy function sensitivity analysis. A remedial measures expert system is also described, but this has not yet been built.

The paper points out that the objective of this effort was to prove the feasibility of the concepts that are described. A particularly interesting comment concerns the weaknesses of commercially available expert system shells available at the time: "... it appears that none are designed to deal effectively with the complexities inherent in this type of application", these having to do primarily with the wide range of data sources and computations required. All parts of the system, including the inference engine, rule bases, and simulation programs were programmed in FORTRAN which, as in the case of Xue, solves the software integration problem in the short term. Unfortunately, here also, no information is given regarding the effectiveness of the system, the simulation programs, the system models used or the size of the networks on which it has been tested.

Fouad, Venkataraman and Davis [Fouad, Venkataraman & Davis 1991] describe an expert system which performs dynamic security assessment on the basis of results obtained using the transient energy function method. The effect of contingencies on system security, and the sensitivity of system security with changing load conditions and parameters are computed using this method. Control parameter settings are recommended which can then be adjusted to maintain the security index within a range of safe values. A particularly interesting part of this effort relates to the application of the expert system to the Northern States Power Company network. This is valuable in that the system clearly attempts to integrate expert system methods, algorithmic stability analysis, a heuristic security index and control parameter correction guidance to a real operating environment. The only reservation one might have here concerns the

generality of the transient energy function for stability evaluation.

Pattern recognition techniques have been applied to power system security, principally by Pang et al. [Pang et al. 1973] and Oh [Oh 1986]. These methods are essentially based on a comparison of the real–time network with a large number of previously–studied cases. Using a strategy that attempts to determine a closest fit, it is then possible to estimate the state of the network (i.e. secure or not) and even perform "what if" analysis on degraded networks. Conceptually, these methods are ideal for porting off–line results into the on–line environment and implementing real–time monitoring. In practice, however, they have been handicapped by very high computing requirements.

Recent developments in artificial neural networks have rekindled interest in this area, particularly with a view of overcoming the weaknesses of numerical pattern–matching techniques. For instance, research has shown that an artificial neural network "... can successfully be trained to capture a very non–linear relationship between controls (such as generator MW and MVAR outputs) and the resulting bus voltages throughout a system" [Avramovic & Fink 1991], both for a base and post–contingency networks, and recall them with very high efficiency. Key challenges in this technique include neural network topology theory and teaching strategies, and many issues remain open. Though the technology is still under development, "a combination of experimental and theoretical knowledge provides an ever–increasing foundation for understanding the behaviour of artificial neural networks" [Avramovic & Fink 1991]. The technology is emerging as a serious candidate in many power system application areas such as power system stabilizers [Malik & Zhang 1992], voltage control [Avramovic & Fink 1991], substation alarm signal filtering, network observability determination, load prediction, power plant model identification, unit commitment, and stability analysis, security and control [El–Sharkawi & Marks 1991].

1.6.6 Probability and dynamic security analysis

Probability methods have also been tried in dynamic security assessment. Some of this work has attempted to deal with the uncertainty related to the nature and location of contingencies. One proposed probability measure is time to insecurity [Wu & Tsai 1983]; another is a so–called *dynamic security measure* of a power system, defined as the integral of the probability density of the state–

structure pair over the dynamic security region, which considers state–independent and state–dependent (i.e. system–response) disturbances [Loparo & Abdel–Malek 1990]. Though the above theories are technically complex, they are usually based on simplifying assumptions which have limited their practical application: In particular, the dynamic security region must be known beforehand and is usually obtained using direct methods.

Leite da Silva, Endrenyi and Wang [Leite da Silva, Endrenyi and Wang 1993] have proposed another approach, based on the considerable work done in relation to power system adequacy: Essentially, they propose an integrated treatment of adequacy and security. This may well provide a useful foundation for future work.

1.7 THE PRESENT THESIS

If one wishes to improve the efficiency and accuracy of off–line dynamic security analysis, where to begin? Getting more experts is clearly impractical: They are already in short supply, and their gestation requires a lengthy individual and corporate commitment. We have seen the relative merits of particular developments using artificial intelligence techniques or expert systems, usually in combination with energy–based methods or simplified algorithmic tools, and some extremely attractive concepts have been described along these lines. However, apart from the early work by Crevier [Crevier 1978], no serious attempt has been made to mechanize traditional dynamic security analysis processes based on tested, commercial time–domain simulation software in the same way that such processes have evolved and are performed in typical operations planning environments.

The present thesis therefore provides a formal overview of dynamic security analysis as practiced in operations planning for the purpose of mechanizing the processes presently realized by experts. A *framework* prototype which mechanizes many of these processes, developed in the course of this work, is described and examples are provided of its performance in a typical operations planning environment. This prototype was also found to be extremely useful in performing research in dynamic security, and some new, fundamental results are presented which 1) show that transfer limits can be estimated quite effectively from at most three stable transient stability simulations and 2) lead to a

simple criterion for the stability evaluation of individual transient stability simulations. Such results could not have been obtained without the help of the framework prototype.

1.7.1 PART 1: A FRAMEWORK FOR MECHANIZING DYNAMIC SECURITY ANALYSIS IN OPERATIONS PLANNING

PART 1 of this thesis introduces the concept of a *framework* [Kaplan 1992] for the purpose of mechanizing processes in dynamic security analysis. A *framework* attempts to combine semantic generality with software generality, where semantic generality refers to the capacity to define specific power transfer limit determination strategies from a large list of possibilities, and software generality represents having the option of using any one of many commercial load–flow and transient/long–term stability software packages.

Though the scope of dynamic security analysis processes is very large, it can be subdivided into essentially two types: 1) traditional dynamic security limit search processes based on the use of simulation software and 2) purely heuristic limit determination processes based on a data base of past off–line simulations. Though both types of processes are described in the present thesis, both being important, the focus here is on limit search processes involving large commercial simulation software. An important consideration in favour of this choice is that *frameworks* for dynamic security analysis based on proven off–line processes and techniques will inevitably gravitate to the on–line environment thanks to cheap, high–speed computing power. Of course, this evolution will heavily depend on the emergence of the appropriate hardware and economics and this, in turn, cannot be predicted with any accuracy.

Chapter 2 provides an introduction to dynamic security analysis in operations planning. This is done principally by means of an example. A survey of the domain knowledge required to mechanize human processes in dynamic security analysis is outlined in this chapter.

Chapter 3 structures the processes of dynamic security analysis, as used in operations planning. A taxonomy of these processes is proposed using a semantic net, enhanced by means of classes, objects, properties and rules. Traditional dynamic security limit search processes are treated in detail. Heuristic pro-

cesses are outlined, and it is shown how the two can be integrated within a unified approach. It is also shown how the semantic net can be viewed as an object–oriented language and how the language syntax is derived from the semantic net. Conditional branching structures and filters for dynamic security analysis are also considered.

Chapter 4 provides a detailed description of the ELISA framework prototype developed within the course of this work. A number of examples of language scripts presently supported by ELISA are given and the results are discussed. Examples are also presented which demonstrate how ELISA is used to perform large–scale sensitivity studies and how this impacts a specific operations planning environment. User interfaces are also discussed.

1.7.2 PART 2: THE ELISA FRAMEWORK AS A RESEARCH TOOL

Because the mechanization of dynamic security limit search processes gives considerable leverage to an expert, accomplishing in hours or days what would normally take weeks or months, ELISA is an extremely attractive tool for research. PART 2 of this thesis therefore describes the use of the ELISA prototype in performing research in two different areas: 1) the search for a performance index which characterizes contingency severity and permits the rapid estimation of power transfer limits, and 2) the search for a simple, mechanizable stability criterion.

Chapter 5 proposes a new stability performance index based on the signal energy of the network voltage transient response. It is shown that this index asymptotically approaches the network transient stability transfer limit for a particular contingency and topology. It is also shown that the index provides information regarding the relative severity of different system operating points in relation to a given contingency, and that the stability limit characterizes the relative severity of different contingencies on the same network. A simple model is derived, explaining this behaviour near the stability limit, and it is shown that at most three stable transient stability simulations are required to find a transient stability limit with high accuracy. The model is shown to be applicable to all normal contingencies, and the results hold notwithstanding the presence of many active, non–linear elements in the network.

Chapter 6 shows how power system stability can be evaluated using a frequency–domain stability criterion. Essentially, this criterion requires performing the Fourier transform of a network's transient voltage response at various monitoring locations: When P is varied and the network goes beyond its stability limit, the angle of the Fourier transform's polar plot fundamentally changes its behaviour, passing from a clockwise to a counterclockwise rotational behaviour about the origin. This is confirmed by results obtained from performing stability–limit searches on the Hydro–Québec system. Used in conjunction with signal energy analysis for determining stability limit proximity, this criterion can be quite useful for mechanized security–limit–determination tools such as ELISA.

Finally, Chapter 7 provides an overview of the present thesis and recommendations are made regarding future work.

1.8 Claim of Originality

The following constitute the principal results of the research presented in this thesis and, to the best of the author's knowledge, represent original contributions to the study of dynamic security analysis:

1. A taxonomy of the principal processes and domain knowledge currently required for off–line dynamic security analysis.

2. A system for presenting said taxonomy based on semantic nets, classes, objects, properties and rules, and which charactererizes the details of these processes by means of appropriate semantics and syntax.

3. An object–oriented language, based on said semantic net, which has the capability to express high–level goals, define lower–level processes and describe the appropriate conditional branching structures and filters required in the achievement of these goals.

4. An approach for implementing said language as a *framework* capable of supporting semantic and software generality. Semantic generality here refers to the capability of describing and executing the processes required by dynamic security analysis (i.e. and identified in the semantic network),

and software generality refers to the concept of accomplishing these processes using any commercially available power flow or transient stability software.

5. The development of a production–grade prototype of this *framework* concept, named ELISA, which mechanizes dynamic security analysis based on transient stability, long–term stability and steady–state post–contingency PV–curve transfer limit search processes.

6. The derivation of an improved formula for approximating the fault impedance required to simulate single–line–to–ground faults in the case of degraded networks.

7. An approach for approximating a power system's impulse and transient response from time–domain transient stability simulations.

8. A discovery that the signal energy of a power system's transient response rises asymptotically towards the transient stability transfer limit.

9. An explanation for this behaviour in terms of a power system's dominant poles and their proximity to the imaginary axis, yielding a formula which permits the estimation of transient stability transfer limits with relatively high accuracy.

10. A strategy for integrating this result in existing transient stability transfer limit search processes, in particular, as a means of obtaining an accurate transfer limit estimate in at most three stable transient stability simulations.

11. An approach for evaluating system stability, based on study of the Fourier transform of the network transient response: If the polar plot of the Fourier transform has a sustained counterclockwise phase behaviour, the power system is unstable.

PART 1

A FRAMEWORK FOR MECHANIZING DYNAMIC SECURITY ANALYSIS IN OPERATIONS PLANNING

CHAPTER 2

INTRODUCTION TO DYNAMIC SECURITY ANALYSIS IN OPERATIONS PLANNING

2.1 INTRODUCTION

Before dynamic security analysis processes can be mechanized, the *domain knowledge* must carefully be surveyed, mapped and characterized in terms of:

a) the goals that are targeted;
b) the related methodologies and their relevant tools;
c) the heuristic knowledge required at every stage;
d) the network components, quantities and characteristics that constitute the variables and parameters of the different processes.

In the preceding chapter, we discussed operations planning dynamic security analysis in general terms, primarily from a high–level industry perspective. In the present chapter, we dissect the many detailed technical issues which must be addressed in order to perform dynamic security analysis. Though numerous references are made to the Hydro–Québec system, these issues are common to all utilities and many of the dynamic security analysis methodologies described herein are in common use throughout the industry. In the next chapter, we will propose a taxonomy of dynamic security analysis processes based on this description which, in turn, provides a foundation for their mechanization.

2.2 THE GOALS OF DYNAMIC SECURITY ANALYSIS

In the previous chapter, we saw that dynamic security analysis encompasses a large class of problems. This includes finding:

1. an appropriate power flow for a given set of injections and loads;
2. the security status (i.e. is the case stable or not?) for a given set of injections and loads relative to some contingency (i.e. fault type, duration and location, followed by a change in topology);
3. the transient or long–term stability transfer limit of a transmission line or corridor with respect to a given contingency;
4. the security limit of a transmission line or corridor with respect to a

given contingency, and

5. the location of the contingency which yields the security limit (i.e. commonly referred to as the worst contingency).

In addition to these, the effect of various network parameters on security limits must often be evaluated: This is referred to as sensitivity analysis. Prime candidates for sensitivity analysis include optimizing a) SVC distribution for a fixed number of SVCs, b) the settings of power system stabilizer systems (i.e. on generator excitation systems) or c) protective relay settings. Additionally, security limit sensitivity to tie line flows, flexible AC transmission system (FACTS) device location [Hingorani 1993] and special protection system (SPS) parameters such as load shedding, generation rejection and shunt inductance switching can also be vitally important. Needless to say, the common denominator in sensitivity analysis is the determination of optimal security limits, and the key to security limit determination is item 3: Transient or long–term stability transfer limit determination. Let us briefly consider why this is so.

Transient or long–term stability transfer limit determination can be viewed as the dividing line between low– and high–level processes in dynamic security analysis. Items 1 and 2 represent the elementary functions required in accomplishing item 3; this, in turn, constitutes the fundamental building block for finding items 4 and 5. For example, security limit determination follows directly from a simple comparison of the different transfer limits obtained for different contingencies at different locations (i.e. the lowest transfer limit is the security limit). Sensitivity analysis represents yet another level of complexity: This requires entering a process in which item 4 becomes the primitive, and this primitive is performed for different values of some network parameter.

From a strategic perspective, we have also seen that transmission capacity limits constitute the lowest level of information which characterizes the intersection of a given topology with system planning criteria (see 1.3.2). It is therefore clear that transient/long–term stability transfer limit determination is the key process in dynamic security analysis. The remainder of this chapter will thus address the prerequisites, assumptions and detailed technical issues surrounding this process. We begin by considering the example of next section.

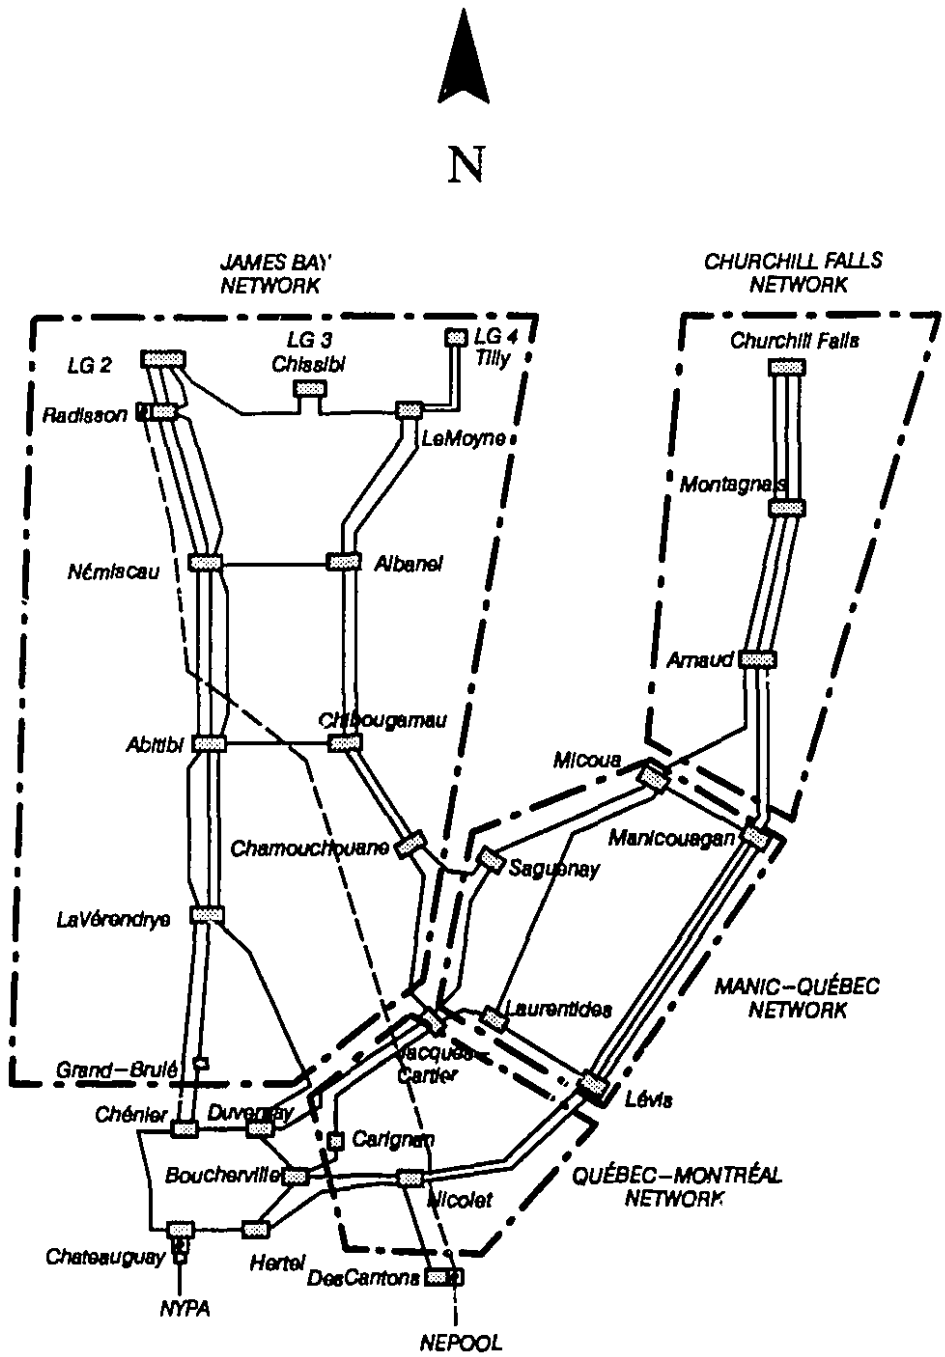

Fig. 2–1. Hydro–Québec's 735 kV transmission system.

2.3 A TYPICAL TRANSIENT STABILITY TRANSFER LIMIT SEARCH

Consider the Hydro–Québec network shown in Fig. 2–1. We will illustrate the steps followed by an operations planner as he or she searches for a transient stability transfer limit using 1) the RP600 commercial power flow software [Gaba, Audette, Guillemette & Lafrance 1993], 2) the ST600 transient/long–term stability software [Valette, Lafrance, Lefebvre & Radakovitz 1987] and 3) a standard 689–bus model of the 1991 Hydro–Québec AC network used for performing production dynamic–security studies. The main Hydro–Québec network consists of some 55 power stations representing a generating capacity of approximately 30 000 MW, 95% of which is hydro–based. The network includes many voltage levels, ranging from 69 to 735 kV. As two–thirds of the total generation is located more than 800 km from the principal load centres in the south, the 735 kV grid consists of more than 10 000 km of lines, and includes 11 static var compensators (SVC) and 9 synchronous condensers (SC) for active voltage support. The main load is centered around the greater Montréal area, primarily supplied by a 735 kV loop consisting of the Chénier, Duvernay, Boucherville, Hertel and Chateauguay switching stations.

2.3.1 THE DEGRADED NETWORK

As we will retrace *all* the steps of the usual process, we begin by degrading some part of the Hydro–Québec network. We choose to focus our attention on the James Bay transmission corridor (Fig. 2–1). First, we remove a long transmission line on the west side of the corridor, that is to say from LG 2 to Chénier 735 kV switching stations, and a single line section between Micoua and Saguenay: We now have an Nd_i degraded network (see Fig. 1–1). We further choose to degrade voltage support by removing two SVCs at Némiscau: This yields the Nd_{ij} degraded network.

2.3.2 THE CHOICE OF CONTINGENCY

Since James Bay transmission capacity has clearly been weakened, a new security limit is necessary. Though we will perform a single transient stability limit search example, an identical contingency must in principle be applied at different locations along the James Bay network to determine the new transfer limit at each one, and the lowest transfer limit obtained is the new security limit.

We choose our contingency C_k to be a single–line–to–ground (SLG) fault of 6–cycle duration. This fault is applied at the Chamouchouane switching sta-

tion, followed by the loss of the line to Jacques Cartier. The simulations are run for a time–length of 600 60–Hz cycles (i.e. 10 seconds) using a winter (i.e. peak) load model.

2.3.3 THE CHOICE OF INJECTION AND LOAD VARIABLES

As described in 1.6.2, to find a transient stability transfer limit on a corridor, one performs transient stability simulations of the same contingency (i.e. fault type and location) for different values of pre–contingency power transfer. This is done by varying the dispatching of the generating stations which inject power into this corridor (i.e. LG 2, LG 3 and LG 4); in every instance, the value of the load at the receiving end is changed accordingly.

2.3.4 POWER FLOW AND VOLTAGE PROFILE

An important step before executing any transient stability simulation consists of performing power flow simulations for the purpose of obtaining appropriate starting conditions from an operations point of view. A flat 1.0 p.u. voltage profile is usually preferred: A uniform voltage profile at rated voltage favours the post–contingency flow of synchronizing power between generating stations when reacting to disturbances. Additionally, all thermal and short–circuit capacities must be respected.

2.3.5 THE TRANSIENT STABILITY LIMIT SEARCH PROCESS

2.3.5.1 THE SEARCH STRATEGY

Figure 2–2 illustrates the results of a transient stability limit search process. We see the voltage in p.u. as a function of time at Duvernay switching station (i.e. near the main load). Each voltage waveform corresponds to the transient stability simulation for a different initial steady–state power transfer on the James Bay transmission system. The cases were generated according to a binary search strategy in the following chronological order:

1) base case
2) base case + 800 MW
3) base case + 400 MW
4) base case + 600 MW
5) base case + 500 MW

The base case sets the combined dispatching of the LG 2, 3 and 4 power stations to 8830 MW. Each of the increments with respect to the base case (i.e. 800,

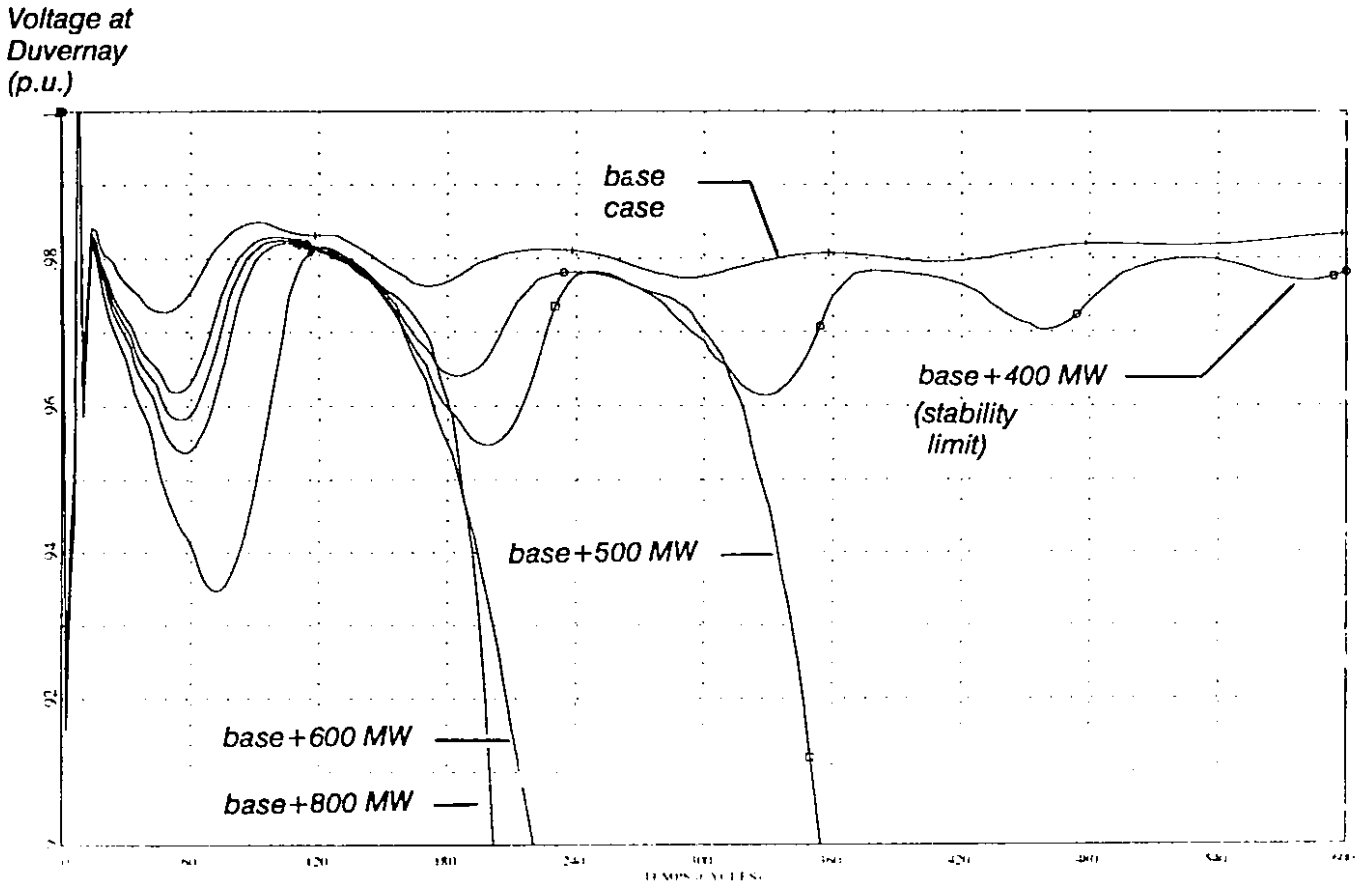

Fig. 2−2. Typical transient stability limit search. Voltage at Duvernay for a SLG fault at Chamouchouane and loss of line to Jacques−Cartier. Each waveform corresponds to a stability simulation for different values of pre−contingency power transfer on the James Bay network.

400, 600 and 500 MW) was distributed on a *pro−rata* basis among the three power−stations. By this, we mean that if the dispatching of the LG 2 power station represents 50% of the combined total in the base case, then 50% of each succeeding increment is allocated to LG 2.

2.3.5.2 ACCEPTABILITY CRITERIA

For a case to be acceptable, it must be stable and its voltage must not exceed emergency limits (i.e. between 0.95 and 1.05 p.u. in the case of Hydro−Québec) after a few seconds. In some instances, other considerations such as tie−line flows or frequency may also be important. Where frequency is an issue, it must remain within certain limits such as 58.5 and 61.3 Hz at every bus. For the purpose of this example, let us suppose that our criterion is to remain within the above−defined voltage limits after 8.33 seconds (i.e. 500 cycles).

2.3.5.3 APPLYING THE CRITERIA

Having established that the base case does indeed meet this criterion, generation at James Bay (i.e. LG 2, LG 3 and LG 4 power stations) is increased by 800 MW which, in turn, injects as much into the James Bay network. Figure 2−2 shows however that this second case (base + 800 MW) is clearly unstable as the voltage collapses before 300 cycles. In fact, this transient voltage collapse phenomenon is characteristic of instability on the Hydro−Québec system: We will return to this in 2.4.9.

Having found a stable and an unstable case within 800 MW of each other, use of the binary search strategy dictates that the user go halfway between the two: The third case tried is thus base + 400 MW and we see that this case also respects our criterion. The stability limit has now been circumscribed between two cases which are 400 MW apart. The fourth *(*base + 600 MW*)* and fifth *(*base + 500 MW*)* cases are both clearly unacceptable, due to transient voltage collapse, and this progressively circumscribes the limit within 200 and 100 MW respectively.

2.3.5.4 TERMINATING THE SEARCH

The process is interrupted at this point as 100 MW corresponds to an acceptable level of precision. By visual inspection of the waveforms, one sees that the third case *(i.e.* base + 400 MW*)* has the highest power transfer which passes the criterion: In practice, this conservative value is taken to be the transient stability transfer limit.

2.4 A DETAILED OVERVIEW OF THE TRANSIENT STABILITY LIMIT SEARCH PROCESS: AN INSIDE LOOK AT OPERATIONS PLANNING

The preceding section has presented a typical transient stability transfer limit search. In this example, the process of determining transient stability transfer limits appears straightforward, but behind this apparent simplicity lie a number of complex issues that have been resolved before embarking in the process, particularly in dealing with uncertainty (i.e. regarding data, load representation, climactic considerations, etc.) and the need for consistency. These issues are so critical that they merit further comment.

Operations planners deal with uncertainty by being highly conservative in their assumptions, methods, approximations and interpretations: For example, a stability limit is defined as the *stable* case with the *highest* power transfer, not as the *unstable* case with the *lowest* power transfer. The consistency requirement may even come as a surprise: How can a process, when performed by different people, using the same tools and data, *not* consistently give the same transfer limits? Such a situation is indeed possible as there exists as yet no standard reference on operations planning philosophy. Even in-house documents are virtually non-existent: Operations planners learn their trade in an apprentice-like environment over a period of years, under the tutoring of senior personnel. Because of the size of the task, approximations are commonplace, and it is difficult to guarantee that the correct ones have been made at the most opportune moment, particularly in a sizeable department where the ultimate control of security limits depends essentially on the experience of senior staff and their "feeling" regarding the network's evolution.

As mentioned before, these concerns are shared by all utilities. An overview of the detailed technical issues which must be addressed in order to deal with these concerns is therefore a primary consideration, and this is the focus of the remainder of this chapter. In order to give substance to these issues, frequent reference is made to the Hydro-Québec system, and many of the latter's practices are described in detail. Because of its numerous HV, EHV and HVDC interconnections, Hydro-Québec maintains close ties to neighbouring NPCC utilities, and many of the dynamic security analysis techniques employed there reflect industry-wide practices. The issues to be discussed are grouped under the following headings:

1. Subdividing a network into transmission corridors
2. Choosing variables to drive the limit search process
3. The dispatching of multiple generating units within a power station
4. Choosing "umbrella" topologies for degraded networks
5. Choosing credible contingency locations
6. Load flow, stability and contingency data
7. Voltage profile and power flow adequacy
8. Single–line–to–ground fault impedance calculations
9. Stability assessment

2.4.1 SUBDIVIDING A NETWORK INTO TRANSMISSION CORRIDORS

2.4.1.1 TRANSFER LIMIT DETERMINATION: THE BASIC PHILOSOPHY

The fundamental approach used for determining the transient stability transfer limit of a single transmission line is as follows: One subjects the line to a credible normal contingency, and this is repeated for different values of line power transfer. Such a contingency can include the loss of the line. The highest value of power transfer compatible with previously–defined acceptability criteria determines the transient stability transfer limit.

2.4.1.2 DEFINING TRANSMISSION CORRIDORS

In principle, a limit must be found for every single transmission line section in a power system. In practice, many lines are in parallel which permits grouping them into *corridors*. Long transmission corridors, consisting of essentially parallel lines, can often be identified where intermediate switching stations tap off relatively little of the power which is being transmitted elsewhere. An alternative definition of corridors is as follows: The power injected by upstream generating stations is divided about equally among all the lines, and power transfer is roughly the same throughout, apart for losses and small loads. Unfortunately, network–specific experience often plays a critical role in one's ability to clarify such fuzzy concepts as "about equally" and "roughly the same".

On this basis, a network can be divided into a certain number of transmission corridors (or transmission networks) for security limit estimation. Of course, this also facilitates the distribution of tasks within a team of operations planners. For example, operations planners at Hydro–Québec usually split the

main 735 kV network into the four corridors shown in Fig. 2–1: 1) James Bay, 2) Churchill Falls, 3) Manic–Québec and 4) Québec–Montréal.

2.4.1.3 CONSTRAINTS ON THE REMAINING NETWORK

A final point. When focusing attention on a single transmission line or corridor, steady–state generation and load are set to peak load conditions in the remaining network. In addition, the latter should remain essentially unaffected by the steady–state power transfer variations imposed on the targeted corridor for transfer limit determination: Otherwise, this will affect the results. The limits obtained by this approach are conservative in general and capable of covering a wide variety of operating conditions. However, light load conditions are also verified and any restrictions due to these are also taken into account.

2.4.2 CHOOSING VARIABLES TO DRIVE THE LIMIT SEARCH PROCESS

To find a transient stability transfer limit, one must subject a corridor to a credible normal contingency for different values of power transfer. A question immediately arises: How is power transfer changed in a corridor while maintaining system conditions unchanged elsewhere? Three different transfer limit determination strategies can be employed to achieve this:

1. The generation–load limit search
2. The load–load limit search
3. The generation–generation limit search

2.4.2.1 THE GENERATION–LOAD LIMIT SEARCH

The method which instinctively arises to control power flow in a corridor is to modify the dispatching of the generating stations which inject power directly into the corridor. However, to correctly simulate system conditions, it is necessary to adjust some other parameter, such as load, at the corridor output, otherwise the power flow equations force the swing bus to absorb the increase (or decrease) in generation. Such a strategy is unsatisfactory in that it leads to limits reflecting the capacity of the targeted corridor and that of the swing bus collector network: When the latter becomes heavily loaded either in the generation or absorption mode, this can artificially restrict the power transfer limit.

The generation–load limit search approach can effectively be exploited whenever sufficient generation and load are present at the corridor extremities. In Fig. 2–1, the James Bay, Churchill Falls and Manic–Québec corridors natu-

rally lend themselves to such a search strategy. In practice, generation and load can be changed at individual buses, groups of buses (i.e. zones), or on the whole network.

2.4.2.2 THE LOAD–LOAD LIMIT SEARCH

In certain cases, it is not possible to adopt a generation–load limit search strategy because of insufficient generation or load at the corridor extremities. An example of this is the Québec–Montréal corridor which suffers from an absence of generation injecting directly into this network (i.e. at its north–east end, at the Laurentides and Lévis switching stations), a sizeable load being found there instead. To use the generation–load approach would require changing generation in the Churchill Falls and/or Manic–Québec corridors, causing the transfers to change in these corridors simultaneously with that of the Québec–Montréal corridor, thereby violating the requirements of 2.4.1.3.

In order to maintain essentially constant power flow along the Churchill Falls and Manic–Québec corridors while changing only that of the Québec–Montréal corridor, the approach used in such a case is to reduce the load at Laurentides and Lévis, and increase that connected to the 735 kV loop around Montréal (i.e. at the south–west end). As in the previous case, load can be changed at individual buses, or at groups of buses (i.e. zones).

2.4.2.3 THE GENERATION–GENERATION LIMIT SEARCH

We have seen that transient stability transfer limit searches can require driving generation at some location(s) and adjusting load at some other location(s), and driving load at some location(s) while adjusting load at some other location(s). Additionally, situations can arise where one wishes to drive the generation at some location(s) while adjusting generation elsewhere. Such conditions can occasionally arise in the management of hydraulic reservoirs, for instance, when one reservoir's level is quite high due to overabundant precipitation while another has lower than average reserves. Or, alternatively, scheduled maintenance may require the partial shutdown of a station: It is consequently dispatched at maximum generation for some period preceding the maintenance in order to forestall the useless discharge of water during repairs. One will also wish to perform such studies in highly degraded networks where generating capacity exceeds that of the transmission system to which it is connected: Certain distributions of power dispatching may have to be avoided.

As a typical example of this, for a given total power transfer on the James Bay network, one may wish to drive (i.e. increase) the generation at LG 3 and adjust (i.e. decrease) generation at LG 2 until a limit has been found for a contingency along the line between LG 3 and LG 2. Once again, in practice, generation can be changed at individual buses, or at groups of buses (i.e. zones).

2.4.2.4 ACCOUNTING FOR LOSSES

Regardless of the combination of network quantities to be driven or adjusted, network losses must also be accounted for whenever the search process causes the losses to vary markedly. In such cases, the swing bus is the only mechanism normally open to the power flow program to compensate for losses. Due to the fact that it is desirable for the swing bus to remain at a relatively constant level throughout a limit search so that its behaviour does not affect stability limit determination, it is best to attempt to compensate for losses in some way.

To this end, operations planners will, for example, increase(decrease) generation by x and increase(decrease) load by $x \times k$, where k is a constant which roughly accounts for losses: Losses are thus asked to *participate* in the increase or decrease of load. The value of k is typically between 0.9 and 0.99. This approach does not completely eliminate swing bus generation changes, but they can be so reduced as to become negligible.

2.4.3 THE DISPATCHING OF MULTIPLE GENERATING UNITS WITHIN A SINGLE POWER STATION

An important corollary to the previous point is that of dispatching multiple generating units within a single power station for transient stability transfer limit determination. The underlying philosophy here is that the dispatching of units, for whatever amount of power requested of a generating station, must reflect the most unfavourable situation from the point of view of transient stability.

Consider a single generating unit. In the event of a fault, the most restrictive condition is that where the net change in angular momentum of the rotating masses upon clearing of the fault is at the maximum possible value (i.e. consistent with gate position): When reconnected to the network, the capacity of the network to provide sufficient electrodynamic braking and resynchronize the

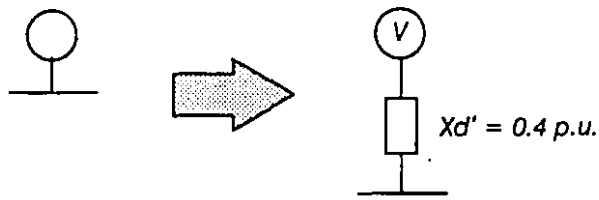

(a) Consider a single 400 MVA generator: $Xd' = 0.4$ p.u. (on machine base).

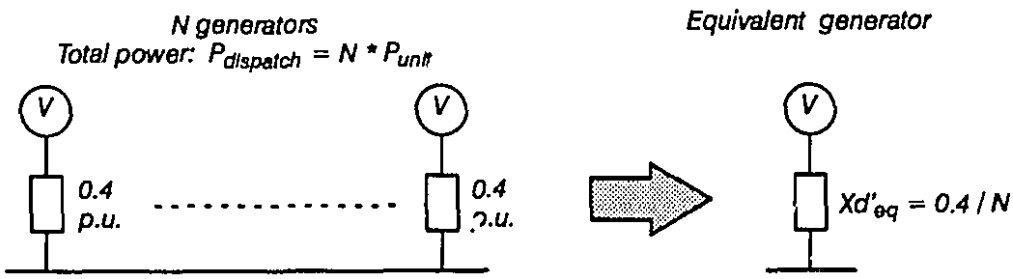

(b) Consider N identical generators generating a total of $P_{dispatch}$:
Since: $N = P_{dispatch} / P_{unit}$
Then: $Xd'_{eq} = 0.4$ p.u. $* P_{unit} / P_{dispatch}$

Xd'_{eq} Impedance Table

No. Units	$P_{dispatch}$ (MVA)	Xd'_{eq} (p.u.)
4	1600	0.10
3	1200	0.13
2	800	0.20
1	400	0.40

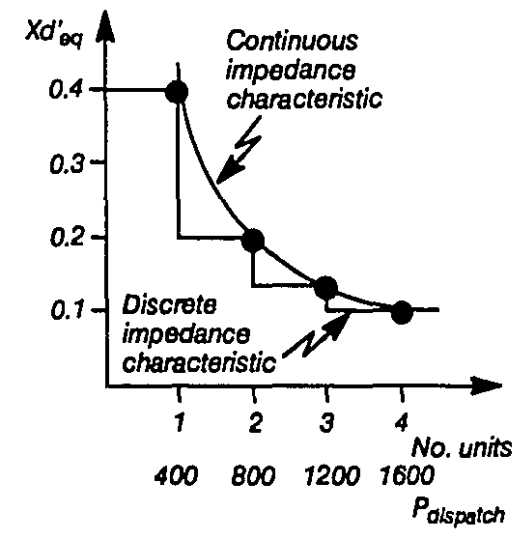

(c) The plot corresponds to the table values for Xd'_{eq}, calculated using the formula given in (b). A continuous impedance characteristic implies a continuously varying N. Xd'_{eq} is given on a single machine base (i.e 400 MVA)

Fig. 2–3. Example calculation of generator transient impedances for transfer limit determination. The same approach holds for all generator impedances.

unit is thus put to its greatest test. This condition represents that of the unit operating at maximum capacity. Translating this in terms of a generating station with multiple units, the most constraining operating condition for a given power dispatching (i.e. from the point of view of transient stability) corresponds to simulating a minimum number of discrete units, with each one operating at nameplate capacity; if required, an additional unit is then provided to pick up the difference in the desired dispatching.

2.4.3.1 GENERATOR CHARACTERISTICS

When redispatching power according to this approach, in the case of multiple–unit hydro generating stations, one must add or delete individual generating units at every generating station used in the limit search process, and adjust the power on at least one additional unit. These manipulations involve making many changes to the data files used as input to power flow and transient or long–term stability packages, and represent a golden opportunity for introducing errors into the limit determination process. Consequently, generating stations composed of many identical units are preferably represented as a single machine of appropriate rating, inertia, impedances, turbine characteristics and controls (i.e. speed governors, excitation systems, power system stabilizers, etc.). If a power station is composed of different machine types or different types of controls, identical units will be grouped together and a small number of equivalent machines will be represented. Where machines have been grouped on the basis of identical controls, the control data of a single unit remains the same for the equivalent machine.

Having performed this transformation, the problem has now been reduced to recalculating inertia constants, machine impedances and charging corresponding to an appropriate number of discrete units. These calculations, followed by subsequent data manipulation, still provide opportunity for error.

A simple variant of this is preferred. One calculates equivalent generator characteristics based on a continuously varying number of units not restricted to discrete integer values. An example of this is shown in Fig. 2–3, where the Xd' generator impedance is specifically considered. Generalizing from this example, for all machine impedances:

$$Z_{dispatch} = \frac{Z_{unit} \times P_{unit}}{P_{dispatch}} \qquad (2-1)$$

where $Z_{dispatch}$ is the impedance corresponding to the new generating station dispatching $P_{dispatch}$, Z_{unit} is the impedance of a single unit, and P_{unit} is the rated power of a single generating unit.

As capacitance increases linearly with increasing numbers of parallel generating units, we use the same approach to consider an equivalent charging as a continuous function of $P_{dispatch}/P_{unit}$, as follows:

$$B_{dispatch} = \frac{B_{unit} \times P_{dispatch}}{P_{unit}} \qquad (2-2)$$

where $B_{dispatch}$ is the new value, B_{unit} is the charging of a single unit and P_{unit} and $P_{dispatch}$ are as before.

As power station inertia also increases linearly with increasing numbers of generating units, we use the same approach to define the equivalent inertia constant as a continuous function of $P_{dispatch}/P_{unit}$:

$$H_{dispatch} = \frac{H_{unit} \times P_{dispatch}}{P_{unit}} \qquad (2-3)$$

where $H_{dispatch}$ is the new value, H_{unit} is the charging of a single unit and P_{unit} and $P_{dispatch}$ are also as before.

These continuously variable characteristics are rigourously equal to those of the true equivalent when power station dispatching is equal to some multiple of unit capacity. However, the values are higher at intermediate values, leading to more conservative results generally. At the end of a study, it is always possible to verify the results by performing a final simulation on the basis of a discrete unit distribution. One way of interpreting these formulae is that they provide an equivalent machine which generates power at maximum gate setting, regardless of the power dispatching required.

2.4.3.2 TRANSFORMER IMPEDANCES

Although individual generators do not always have dedicated unit–connected transformers, particularly in large multi–unit hydro power stations, transformer banks are required to integrate generation to the main grid. Consequently, a per unit transformer impedance is always directly associated to an individual generator. For the reasons previously outlined, transformers are treated in the same way as generators, and tranformer impedance and charging are computed using equations (2–1) and (2–2). However, contrary to generator data which is found in the stability software input file, transformer data are usually integrated to the power flow input data .

2.4.4 CHOOSING "UMBRELLA" TOPOLOGIES FOR DEGRADED NETWORKS

The issue of identifying appropriate degraded networks can be resolved in either of two ways:

i) The network is systematically degraded and tested according to the scheme described in 1.6.2; or
ii) The problem is scaled down in some way by means of network–specific considerations and acceptable approximations.

The first method is highly exhaustive and yields a transfer limit for every conceivable degraded topology. However, as we have already seen in Chapter 1, this approach requires treating a very high number of individual topologies which will overwhelm, in practice, the resources of any utility. Operations planners will consequently attempt to reduce the size of the problem using a two–pronged approach (Fig. 1–1):

1. A small set of (Nd_i) "umbrella" topologies (i.e. focusing only on transmission line configurations) is determined in order to *cover* all credible degraded topologies.
2. A small set of (Nd_{ij}) voltage support configurations (i.e. including SVCs, SCs, etc.) is also determined which will permit security limits to be *deduced* for all other credible voltage support configurations.

Consider the James Bay network of Figs. 2–1 and 2–4(a). This system can essentially be regarded as a five–line point–to–point network between the

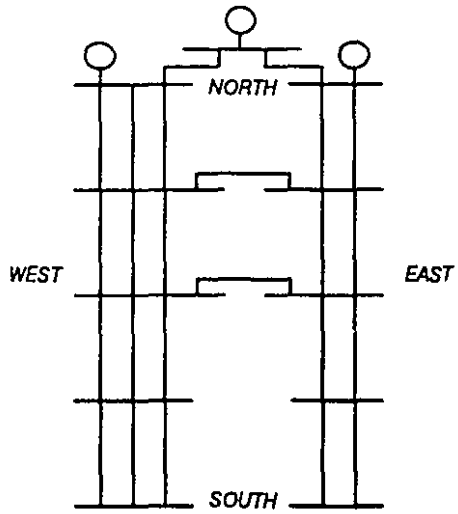

(a) Complete 5–line 735 kV James Bay network

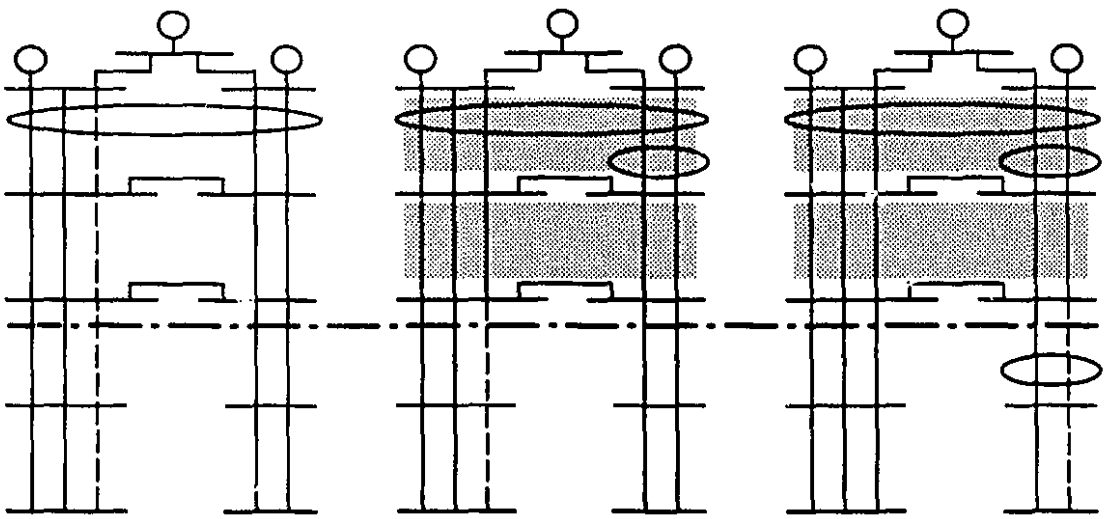

4–line degraded 735 kV James Bay networks: 3 umbrella topologies

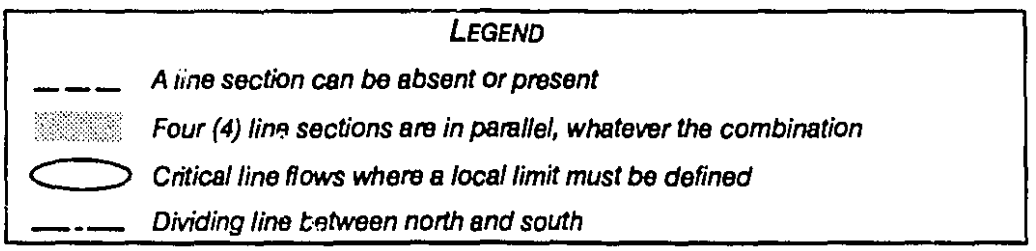

Fig. 2–4. Choosing umbrella topologies for degraded networks. In this example, we see the three umbrella topologies for 4–line James Bay degraded networks.

large James Bay generating complex (i.e. more than 10,000 MW) and the south part of the system: Relatively little load is connected to intermediate switching stations. Operations planners at Hydro–Québec usually work with three families of degraded topologies on this part of the network: The five–, four– and three–line topologies. These families are further subdivided into specific umbrella configurations. Take, for example, the four–line James Bay umbrella topologies shown in Fig. 2–4. In all cases, four complete lines must travel from north to south. A single limit is determined for each one of these umbrella topologies and these three limits, in turn, cover *all* four–line topologies. Of course, some limits will underestimate the capacity of certain configurations, but this is considered an acceptable compromise in dealing with uncertainty. Let us briefly consider these umbrella topologies.

Figure 2–4 (b) shows the basic four–line topology composed of two complete lines along the east and west sides of the corridor. Whether one or four series line sections are absent on the west side, a single limit is chosen to cover all combinations of this particular topology. One clearly sees that the loss of a single line section is considerably less severe than that of four lines in series. Nevertheless, this is considered an acceptable umbrella topology.

Figure 2–4 (c) shows a more complex topology where the main focus is on the north part of the corridor. Here, because of the tie lines which link east and west 735 kV switching stations, all the shaded–in line sections are considered as being in parallel. Any combination of four line sections, for each one of the two northern stages (i.e. two sequences of line sections are shaded–in) is considered an appropriate four–line approximation. In addition to this, any combination of one or two series line sections on the south–west side may be absent. Here again, a single topology covers all the possible combinations.

Finally, Fig. 2–4 (d) shows a variant of the preceding case. The north part of the corridor is treated as before. However, one or two series line sections may be absent on the south–east side. As before, a single limit covers all possible variations of this particular umbrella configuration.

Such an approach has four important consequences:

1. Security limits will generally be conservative, particularly in the case of lightly degraded topologies.
2. A global transfer limit is often insufficient: local limits must also be established for certain critical locations in order to provide adequate security coverage (for example, when a single line section can be found in series with two parallel lines).
3. Such local limits, in turn, constitute important parameters of the security control strategies provided to the system operator.
4. For networks depending heavily on active voltage support (i.e. SVCs, SCs, static excitations systems, etc.), a sensitivity analysis of the impact of this factor on security limits is essential.

Experiential knowledge and creativity are essential in structuring umbrella networks. Often, at a very early stage, a certain number of simulations are defined and performed for the sole reason of testing basic hypotheses (i.e. concerning the network's behaviour) and designing appropriate configurations. Operations planners often describe these tasks as being among the most creative and rewarding.

2.4.5 CHOOSING CREDIBLE CONTINGENCY LOCATIONS

For networks with relatively little installed voltage support equipment (i.e. SVCs, synchronous condensers, etc.), experienced system planners often begin their search for the worst normal continency in the vicinity of the largest generating station. For networks incorporating such technologies, this may constitute a good starting point, but the worst contingency may require performing transient stability transfer limit simulations at a large number of locations.

In the case of degraded networks, a good starting strategy is to begin transfer limit searches in the vicinity of the most weakened parts of the network, that is, where the most elements have been withdrawn from service, or near heavily–loaded lines, particularly tie lines. Apart from these simple guidelines, network–specific experience is often essential in circumscribing credible contingency locations for worst contingency determination. Otherwise, a tentative list must initially be drawn followed by the identification of additional locations on the basis of ongoing study results.

2.4.6 LOAD FLOW, STABILITY AND CONTINGENCY DATA

To perform transient or long–term stability transfer limit studies, one requires power flow and stability data. Power flow simulation describes steady–state network conditions and typically requires such information as real and reactive unit ratings, loads, network topology, swing bus location and passive element characteristics such as impedances and charging of lines and transformers. Where HVDC and FACTS systems must be included, additional characteristics such as firing and extinction angles, and commutating resistances are also important. Stability data are required for the large–scale simulation of dynamic network events: These include turbine, rotating machine, SVC, SC, FACTS device models and associated control system characteristics (i.e. speed governor, excitation system, power system stabilizer, control strategies, etc.), including appropriate HVDC controls. This can also include relay, transformer tap changer, shunt reactor disconnect control (i.e. for long–term voltage stability simulations) and contingency data. Power flow and stability data are typically stored in ASCII files as card images which constitute the input to FORTRAN–based simulation software.

Notwithstanding the power and sophistication of existing simulation tools, the availability of precise data remains problematic for any of several reasons:

1. The age of many power, transformer and switching stations, and their associated protections, and the loss or misplacement of original documents over time;
2. The difficulty to match available models in the simulation software to control systems in the field, particularly for older systems;
3. The unavailability of certain models for newer control systems or new generations of equipment in the simulation software;
4. Line and transformer data which may have been passed on from one generation of planners to another, without comprehensive verification;
5. Load data which may well merit more study, particularly from the point of view of its transient response characteristics.

In many cases, reliable data cannot be obtained by any other means than field tests. However, because of the difficulty of justifying their cost in relation

to short–term economic gain, they are rarely performed. Rotating machine or transformer data are frequently obtained by using "known" data for units of similar size, age and type, such as that found in [Anderson & Fouad 1977].

In general, power system data is *calibrated* using the network's voltage response to known contingencies registered on sequence–of–events recorders found at different network locations. For the operations planner, the proper reproduction of such phenomena with his own tools has precedence over the marginal satisfaction of using what may be construed as more precise data, particularly in the case of control systems.

2.4.7 VOLTAGE PROFILE AND POWER FLOW ADEQUACY

It is a fundamental tenet of power system operations that voltage profile be maintained at 1.0 p.u. voltage: If this is not the case, operators are instructed to perform corrective action. A uniform pre–contingency voltage profile at rated (i.e. 1.0 p.u.) voltage favours the post–contingency flow of synchronizing power between generating stations and hastens system recovery. It follows that, for a given power flow, variations in voltage profile can affect the severity of transient stability simulations, particularly where voltage is low. Additionally, it is important that limits be determined from the same starting voltage profile so that transmission capacities for different degraded networks and contingencies be comparable on an equal footing.

Of course, the necessity of maintaining a 1.0 p.u. flat voltage profile extends to the transfer limit process itself: The pre–contingency network must always be assumed to have such a profile. However, when a team of experts work simultaneously at determining limits, it is difficult to guard against discrepancies, particularly when standard (i.e. as opposed to optimal, or OPF) power flow software is used. The manual adjustment of voltage profile requires entering a lengthy iterative process of redispatching generation, and adding or withdrawing available reactive power components at different locations until the objective is attained. Where the work load is high and time is limited, precision is bound to be sacrificed.

A simple approach to maintaining a flat voltage profile is to transform every bus on the main EHV grid into a PV bus with a "fictitious" 1.0 p.u. voltage generator of zero real power capacity. Having performed the power flow simulation,

these buses are converted to PQ buses before the result is read by the stability simulation software: This transforms the fictitious reactive power sources into shunt impedances. Conceptually, this amounts to using continuous rather than discrete shunt impedance values for steady–state reactive power support.

An underlying assumption in this approach is that sufficient reactive power reserves be resident in the network, whether in the form of shunt inductors, shunt capacitors, SCs or SVCs. When a limit has been found, one need only verify that reactive power flows generated by the continuously variable impedances can be met by appropriate arrangement of the available equipment such that voltage is everywhere equal to or higher than 1.0 p.u.

An important consequence of this approach is that it is no longer necessary to loop through an iterative procedure to obtain an adequate power flow simulation while engaged within a transient stability transfer limit search process. The application of these techniques results in an appropriate power flow simulation in a single step, though subject to verification at the end of a limit search.

2.4.7.1 FICTITIOUS GENERATORS AND PASSIVE SHUNT ELEMENTS

The approach is as follows: Fictitious generators are created on all EHV buses. For example, on the Hydro–Québec network, this is done on every 735 kV bus. Though a 1.0 p.u. voltage is preferred, other values can be used according to the study objective. Lower voltage profiles will result in more conservative (i.e. lower) limits.

Real power is set to zero at all of the fictitious generators. The reactive power range of each one reflects the passive reactive power reserves available at each individual bus. For example, if four 330 MVAR shunt reactor banks are available at a given location, the allowed reactive power range will be entered as 0 to −1320 MVARs. The fictitious generators replace the available passive shunt reserves: Consequently, the latter are removed from the power flow simulation. In the stability data, the load characteristic of these reactive power sources must be entered as varying as a function of the square of voltage.

2.4.7.2 STATIC VAR COMPENSATORS AND SYNCHRONOUS CONDENSERS

Many networks include active reactive power voltage support components such as static var compensators (SVCs) or synchronous condensers (SCs). From an operations perspective, these contribute to:

1. maintaining a flat voltage profile;
2. attenuating system overvoltages;
3. improving power system stability.

Many SVCs on the Hydro–Québec network have a reactive power range of –100 to +300 MVARs. In order to favour functions 2 and 3 above, the system operator tries to adjust steady–state system conditions so that SVCs are functioning very close to 0 MVAR. From an operations planning perspective, the worst operating condition is arbitrarily taken to be that when *all* SVCs are simultaneously generating +100 MVARs to support voltage. Consequently, the steady–state power flow preceding a transient stability simulation reflects this situation. Synchronous condensers are treated in the same way.

To arrive at this, each SVC is modelled as a PV bus in the power flow simulation with a reactive power range constrained to a minimum of +100 MVARs and a maximum of +100 MVARs. This forces the PV bus to generate this value in steady–state, allowing bus voltage to vary. The appropriate SVC controls are included in the stability data. As before, the same holds for SCs.

2.4.8 SINGLE–LINE–TO–GROUND FAULT IMPEDANCE CALCULATIONS

Where faults other than three–phase must be simulated, the standard approach is to consider a power system from a sequence network point of view. This typically involves introducing an impedance at the point of fault, for the duration of the fault, which represents appropriately connected negative and zero sequence networks [Westinghouse 1950]. Network sequence impedances are obtained by performing short–circuit studies of the network employing appropriate short–circuit software. To simulate clearing of the fault, the impedance is removed. Negative and zero sequence contributions are ignored in the subsequent time–domain simulation.

In the case of single–line–to–ground (SLG) fault simulations, the nega-

tive and zero sequence networks are simply added in series with the detailed positive sequence network at the point of fault, as illustrated in Fig. 2–5. In principle, a short–circuit study should be performed for every change brought about to network topology to obtain appropriate negative and zero sequence impedances at the point of fault. If one is to be rigourous, this includes changes to power dispatching (as this involves adjusting generator and transformer impedances, see 2.4.3) which also impacts shunt reactor distribution (i.e. required to maintain a 1.0 p.u. voltage profile, see 2.4.7). The issue of performing short–circuit studies at every step of a transfer limit search can render an already lengthy process considerably more unwieldy.

The following assumption is thus made: Whenever network topology or power dispatching is changed, though the absolute values of the positive, negative and zero sequence impedances also change, it is reasonable to suppose that their relative proportions change very little. Expressing this in different terms and referring to Fig. 2–5, this amounts to assuming that V_f / V remains essentially constant for such changes. When the pre–fault open circuit voltage at the point of fault (i.e. V) is 1.0 p.u., the on–fault voltage drop across the negative and zero sequence fault impedances (i.e. V_f) is thus taken to be constant.

2.4.8.1 OVERALL APPROACH

The approach used for updating SLG fault impedances in the stability transfer limit search process is essentially as follows:

1. A comprehensive short–circuit study of single–line–to–ground faults is performed on the main EHV grid of the complete (i.e. noble) network for a given calendar year. This is done for a network having a flat 1.0 p.u. voltage profile at peak load;
2. A table including the impedance, Z_f, of the combined negative and zero network impedances (i.e. in series) for the complete network, and the associated desired voltage drop V_f is obtained as a function of fault location;
3. When network topology or power dispatching is changed, one obtains a new fault impedance from this information. This fault impedance is then included in the stability software contingency data.

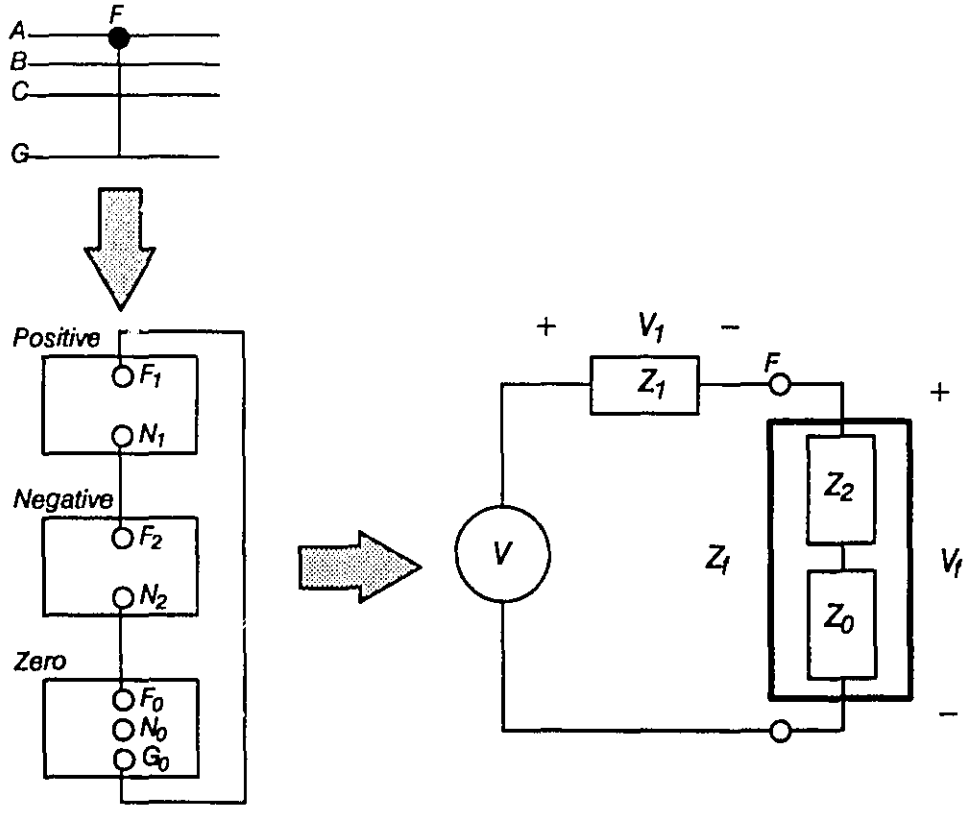

(a) *Connection of sequence networks for SLG fault simulation [from Westinghouse 1950]. V_f / V is considered constant for all degraded topologies.*

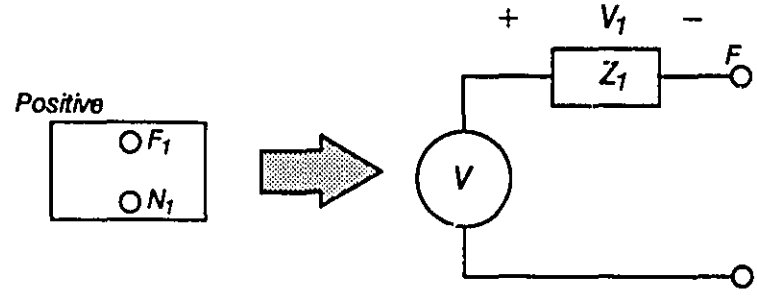

(b) *For calculating fault impedances, the positive sequence network is represented by its pre−fault Thévenin equivalent, looking into the network at F.*

Fig. 2−5. Sequence networks used for SLG fault simulation and impedance determination.

One issue remains: How is the updated fault impedance obtained from the table data as indicated in item 3? This is done as follows:

1. A very short transient stability simulation is performed, typically of one−cycle duration;
2. The combined negative and zero sequence fault impedance Z_f found in the table (i.e. for the appropriate location) is used to simulate the SLG fault in the short simulation. In other words, this is considered a test value;
3. The time step used is very short, on the order or 0.1 or 0.2 cycles;
4. From the simulation results, the positive sequence voltage V_f across the test fault impedance Z_f is obtained at the very first time step after t=0. The time step is kept short so that active system components have not had time to react to the fault;
5. Using this value for V_f, and knowing the test fault impedance Z_f and the desired fault voltage found in the table V_{fd}, one can now obtain the corrected fault impedance Z_{fd} using an appropriate formula.

At Hydro−Québec, an empirical formula has traditionally been used to determine Z_{fd}. In the course of this work, a new formula has been derived which gives better results and which we will now present.

2.4.8.2 COMPUTATION OF SLG FAULT IMPEDANCES

An SLG fault is simulated by connecting the sequence networks in series at the point of fault, as illustrated in Fig. 2−5. The voltage V is the pre−fault (i.e. at $t = 0^-$) equivalent voltage behind the positive sequence network, as seen at the bus in the network where the fault is to be applied. Considering the network as a Thévenin equivalent, V is simply the steady−state pre−fault voltage at this bus, that is before insertion of the fault impedance. Though V usually equals 1.0 p.u., we will consider the more general case.

Let us first consider the short test simulation. We simulate an SLG fault to the network by inserting the test fault impedance Z_f between the bus to be faulted and ground (i.e. at $t = 0$). A voltage V_f is generated across the fault impedance at $t = 0^+$. The following ratio can be written for the voltages across the positive sequence and fault impedances:

$$\frac{V_f}{V_1} = \frac{Z_f I_f}{Z_1 I_f} = \frac{Z_f}{Z_1} \qquad (2-4)$$

where all I_f is the fault current.

The voltage across the positive sequence network V_1 is not explicitly known. It can be obtained from:

$$V_1 = V - V_f \qquad (2-5)$$

Substituting (2–5) for V_1 in (2–4) and reversing the order of the equation,

$$\frac{Z_f}{Z_1} = \frac{V_f}{V - V_f} \qquad (2-6)$$

which expresses the ratio of the fault to positive sequence impedances in terms of pre–fault and on–fault voltages at the point of fault, respectively V and V_f. V_f is obtained from the short stability simulation described previously using the test impedance Z_f.

We now wish to determine the fault impedance Z_{fd} which will yield the desired voltage V_{fd} at the point of fault (i.e. from short–circuit tables). Assuming an SLG fault is applied using the Z_{fd} fault impedance, we rewrite equation (2–6) as follows:

$$\frac{Z_{fd}}{Z_1} = \frac{V_{fd}}{V - V_{fd}} \qquad (2-7)$$

Taking now the ratio of equations (2–6) and (2–7), we have:

$$\frac{\dfrac{Z_{fd}}{Z_1}}{\dfrac{Z_f}{Z_1}} = \frac{\dfrac{V_{fd}}{V - V_{fd}}}{\dfrac{V_f}{V - V_f}} \qquad (2-8)$$

This, in turn, can be reduced to:

$$Z_{fd} = Z_f \times \frac{V_{fd}}{V_f} \times \frac{V - V_f}{V - V_{fd}} \qquad (2-9)$$

which is the desired result. To summarize, Z_f and V_{fd} are obtained from the table of previously computed short-circuit results, V is the open circuit pre-fault voltage at the desired fault location, V_f is the voltage obtained from a test stability simulation using the test impedance Z_f, and Z_{fd} is the desired SLG fault impedance.

2.4.9 STABILITY ASSESSMENT

2.4.9.1 A DISCUSSION OF HYDRO-QUÉBEC PRACTICES

Power system instability is intimately associated with the loss of synchronism, and machine frequency and angle are the quantities most commonly associated with stability assessment. Energy methods focus on the transient energy behaviour of rotating masses which amounts to studying machine frequencies and angles from a different perspective. From a physical point of view, these provide a sound theoretical underpinning for understanding a wide variety of phenomena. Nevertheless, operations planners at Hydro-Québec and other utilities often demonstrate a marked preference for analyzing stability on the basis of voltage behaviour. Let us attempt to understand why.

First, RMS voltage is the primary electrical quantity associated with a power system's capacity to transmit power. On an extensive EHV network such as the Hydro-Québec 735 kV system, voltage control is also an important aspect of system security: One must protect the system against temporary overvoltages in steady-state, and support voltage when large-scale disturbances occur. A good understanding of voltage behaviour under a variety of circumstances is fundamental to one's ability to accomplish this.

Second, when disturbances occur, sequence-of-events recorders located at different stations record the time-dependent voltage variations which result: Voltage has the advantage of being easily measurable anywhere on the network. Additionally, where frequency information is available, it is almost always derived from voltage measurements. Because of this, operations planners have

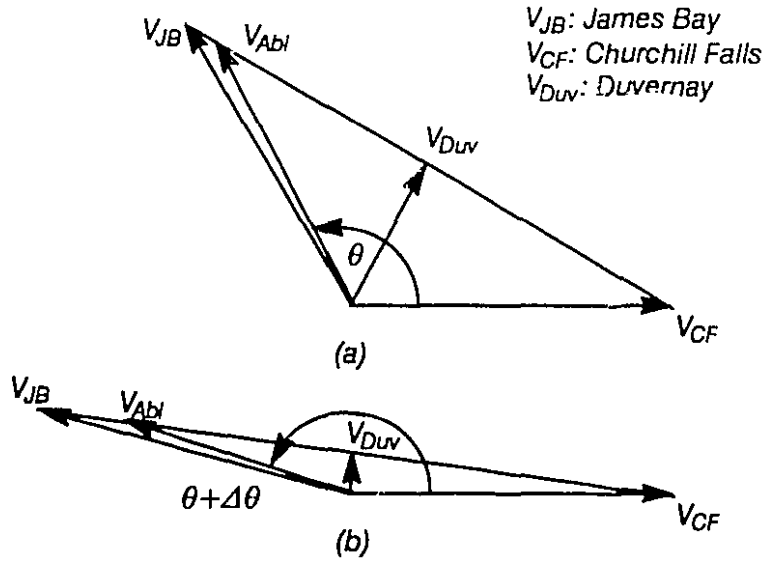

Fig. 2–6. As the V_{BJ} and V_{CF} voltage phasors move away from each other by $\Delta\theta$, the voltage phasors tend to collapse at intermediate points such as at V_{Duv} (V_{BJ} and V_{CF} are assumed to remain constant).

learned to interpret a variety of network phenomena through study of the system voltage waveforms, and simulated system responses are calibrated on this basis.

Though traditional stability assessment proceeds on the basis of frequency and angle analysis, is it possible that voltage also yields sufficient information to assess power system stability? The answer is that stability and instability are *system states*. Instability of a single quantity implicates instability of the whole. Thus if a power system is unstable, all network quantities must reflect instability, regardless of their location in the network, though some quantities may demonstrate more *sensitivity* to unstable behaviour than others. From another perspective, one can state that voltage at any bus is dependent on the network equations which include the contributions of all generators (i.e. voltages, angles and frequencies) and voltage support equipment. System voltage fluctuations thus invariably integrate the changing conditions occurring *everywhere* in the system.

For example, when the Hydro–Québec network begins to lose synchronism, except for stations near generator buses, voltage tends to collapse everywhere, and most severely at intermediate points between the main generating

centers. This is observed from the voltage waveforms of Fig. 2–2 where all unstable cases (i.e. base+800 MW, base+600 MW and base+500 MW) exhibit voltage collapse. This behaviour can be visualized in terms of the simple model of Fig. 2–6: When voltage phasors representing the two main generating zones (i.e. V_{JB} for James Bay, V_{CF} for the combined Churchill Falls and Manicouagan complexes) begin to move away from each other, the voltage drops rapidly at all intermediate points (such as at Duvernay). This shows that voltage–based analysis can provide a meaningful alternative to more traditional approaches based on machine angle and frequency.

2.4.9.2 STABILITY CRITERIA

When studying a network for the first time, all key variables (i.e. frequency, angle, voltage, etc.) should ideally be monitored, for example, at the main generating stations and near important system loads. When an operations planner has gained familiarity with a particular network, the choice of one variable over another, and network location, often rests on the weight of historic precedent, based on past problems and their respective solutions: In other words, operating experience and network simulation experience are both extremely valuable in determining the variables most sensitive to instability.

Stability analysis can also depend on other factors. For example, when optimizing some particular relay setting, the focus will naturally gravitate to those network quantities most directly concerned, and other variables will be accorded secondary importance if time constraints are too great. In addition, on an extensive EHV system, voltage control will always represent an important issue, regardless of whether the focus is steady–state or transient phenomena, and this must be added to other concerns. Such considerations dictate *a priori* that power system stability analysis is as much an art as it is a science.

Typically, for a transient stability simulation to be acceptable, certain network quantites (chosen by the operations planner) must recover, within some period of time (also chosen at the planner's discretion), to pre–defined emergency limits. Such an approach constitutes an implicit stability criterion in that unstable simulations are quite incapable of meeting such conditions. System stability criteria are thus subsumed within appropriate acceptability criteria, defined as firm boundaries to be respected by the quantities concerned. Though

operations planners will provide rigourous technical definitions of the term "emergency limits", even these must be considered specific to a given system.

For example, on the Hydro–Québec system, the voltage behaviour at all switching stations near the main load is a prime concern and the voltage waveforms of any one bus on the relatively short 735 kV lines forming a loop around Montréal are carefully examined. In Fig. 2–2, our criterion thus required that the voltage at Duvernay switching station be greater than or equal to 0.95 p.u. after 8.33 seconds (i.e. 500 cycles). We will return to these issues in Chapter 9.

2.5 SUMMARY

In this chapter, we have attempted to provide an overview of dynamic security analysis in operations planning. This was accomplished by focusing on a key dynamic security analysis process, the transient/long–term stability transfer limit process. Because transient or long–term stability transfer limits represent the lowest level of information which characterizes the intersection of topology and system planning criteria, the underlying processes constitute the foundations of operations planning goals and methodologies. This is therefore a logical starting point for reviewing the highly conservative philosophy, methods, and technical considerations which are typical of operations planning. Though many of the examples relate to one specific utility, the issues presented in this chapter – and their subsequent treatment – are representative of similar concerns and approaches in the industry.

Transient and long–term stability transfer limit determination are simulation–intensive, iterative processes which require using different software tools at differents steps of every iteration, and performing high–level choices and data manipulation at every step. Two underlying concerns are of particular significance: Uncertainty regarding operating conditions (i.e. climactic, topological, etc.), and consistency in limit determination. To address these concerns, appropriately conservative and simplifying assumptions are made when performing such tasks as 1) subdividing the network, 2) choosing appropriate variables to drive the limit search process, 3) dispatching generating units, 4) choosing umbrella topologies for degraded networks, 5) choosing credible contingency locations, 6) determining data, 7) establishing voltage profile and load flow adequacy, 8) calculating fault impedances, and 9) assessing stability.

However, the primary motivation for describing dynamic security analysis practices in operations planning is this: Before such processes can be mechanized, they must first be identified, their assumptions, goals and parts described, and the various relationships defined and characterized. This step, accomplished in the present chapter, is a prerequisite to structuring an appropriate process taxonomy which will facilitate their mechanization. We address the task of building such a taxonomy in the next chapter.

CHAPTER 3

A TAXONOMY OF DYNAMIC SECURITY ANALYSIS IN OPERATIONS PLANNING

3.1 INTRODUCTION

In chapter 2, we saw that dynamic security analysis is a class of problems where individual problems can be solved by the use of distinct, high–level processes. Depending on the specific problem to be addressed, a high–level process is built up from lower–level processes which may consist, in turn, of yet more elementary processes. In this respect, we saw that the transient stability transfer limit search process is fundamentally different from all other processes: It constitutes the lowest level of information which characterizes the intersection of topology and system planning criteria (see 1.3.2), and thereby represents the dividing line between high– and low–level processes.

As we have observed many times, the objective of this part of the thesis is the mechanization of human processes in dynamic security analysis, with a particular focus on operations planning. From a knowledge–engineering point of view, one may see the previous chapter as describing the domain knowledge required to perform dynamic security analysis. A key feature to structuring knowledge in dynamic security analysis resides in the description of the *processes* performed by an expert: Within each process, specific technical issues are then addressed which depend, in turn, on the *goal* targeted by the process.

This leads us to the issue of knowledge representation. In the present chapter, this is the primary concern: We will structure the domain knowledge in terms of an appropriate knowledge paradigm. The resulting dissection of dynamic security analysis into its component parts, appropriately classified, constitutes a taxonomy which can only assist in mechanizing the processes of dynamic security analysis. Such a taxonomy also has an added value: In circumscribing what is known about a domain, one is in a better position to identify those areas where further research is warranted.

To capture the different kinds of knowledge needed to describe dynamic security analysis, the knowledge paradigm must satisfy certain prerequisites. First, it must possess sufficient *richness* to include the goals, processes, variables, parameters and relationships that characterize dynamic security analysis: This is referred to as representational adequacy [Lim & Cherkassky 1992]. A corollary to this is query flexibility: This relates to the ease of formulating queries in terms of the constituent parts of the knowledge representation. Of course, when looking towards the future, the paradigm must also be sufficiently flexible to allow for the integration of new knowledge as it emerges. Last but not least, the representation must also be *useful*: the exercise of structuring knowledge should provide a deeper insight, facilitating our objective of ultimately mechanizing dynamic security analysis processes. Let us now consider the different knowledge–representation paradigms.

3.2.1 THE KNOWLEDGE REPRESENTATION PARADIGMS

The main paradigms presently used to represent knowledge are:

1. Semantic nets
2. Frames and instances, or classes and objects
3. Rules

Many good references cover each of these with varying degrees of detail. An excellent general reference is Waterman [Waterman 1986] which discusses their use in the context of expert systems. Genesereth and Nilsson [Genesereth & Nilsson 1988] also considers each one but from a more rigourous mathematical standpoint. Dillon and Laughton [Dillon & Laughton 1990] provides an introduction and examples of these paradigms as they are presently employed in the field of power systems.

The semantic net (or semantic network) paradigm is a method of representing knowledge which has been found particularly suitable in structuring languages. It represents a form of generalized flow chart, consisting of nodes (representing concepts, things, physical quantities, or anything else which may be appropriate in the context) linked together by directed arcs (which define the relations between the nodes). The arcs are also viewed as establishing a prop-

erty inheritance hierarchy in the net.

Frames are another useful knowledge–representation paradigm where attributes about objects or concepts are stored in slots. Database concepts are useful in understanding frames: At an elementary level, a frame can be viewed as a record format which serves to identify data elements (i.e. the slots) in the database. An instance of a frame can thus be seen as a specific record containing data. Frames include information concerning their relationships with other frames, for instance in a hierarchy, and the concept of inheritance is allowed, thus enabling lower–level elements to possess the characteristics of higher–level elements. Frames can also include information about how to use the frame, or what one can expect to happen next when certain slot values change, or what should be done in the event that certain expectations are not confirmed (adapted from [Minsky 1975]). As we can see, data–driven or slot–driven procedures, often called demons, give frame–based systems considerable leverage in accelerating searches and other processes.

Classes and objects constitute a paradigm conceptually very close to that of frames and their related instances. In this case, the equivalent of a frame is a class, and specific instances are called objects. Class and object characteristics are stored as properties or attributes rather than in slots. Classes and objects can also be organized as hierarchies in which lower–level elements share the properties of higher–level elements through inheritance. In principle, objects communicate with each other by means of messages, not procedures, though messages are data–driven as in the case of frames, and practical implementations prefer use of the term demon [Neuron Data 1991]. It appears that personal preference will dictate the use of frame/instance or class/object vocabulary: Some authors consider them to be virtually indistinguishable [Parsaye & Chignell 1988].

Rules are yet another form of knowledge representation which exploits conceptual or factual relations between concepts, things or physical quantities when these relations can be structured in the form of an IF ... THEN ... statement. Empirical or heuristic knowledge is often stored in this form, particularly when deep knowledge (i.e. a fundamental understanding in the form of a model, an algebraic or differential equation, etc.) is unavailable or inappropriate. For

example, we have seen that stability assessment is often based on heuristic rules which reflect one's knowledge of network–specific quantities most sensitive to instability (see 2.4.9).

Three conclusions arise from this overview:

1. Semantic networks are particularly appropriate for structuring logical flows. Processes are consequently a natural target for this particular representation;
2. Classes, objects and their characteristic properties are a powerful means of *categorizing* things, quantities and concepts, for example when these are found at the nodes of semantic networks;
3. Rules can be used to establish relations between classes, objects, properties and higher–level concepts. Rules can also be represented within semantic networks.

3.2.2 Representing the Knowledge of Dynamic Security Analysis

We have seen that dynamic security analysis is characterized by processes: Semantic nets are thus a good choice from this perspective. Higher–level processes are constructed by means of rules: This means that the rule paradigm cannot be ignored. Additionally, process details (i.e. the concepts, variables, quantities, identifiers, etc.) cannot all be placed on an even footing, and hierarchical relationships can often be identified: This implies that the class–object paradigm can also be useful.

These considerations have led to proposing a *combined* semantic net, class–object and rule paradigm in order to cover the full range of knowledge types which characterize dynamic security analysis:

a) The logical flow inherent to semantic networks provides the underlying skeleton required for the representation of elementary *processes*;
b) Rules permit the *combination* of processes on the basis of intermediate results, opening the door to the construction of more complex processes.
c) Classes, objects and their related properties add *conceptual depth* to the representation, particularly as a means of characterizing the

hierachical relationships between certain nodes in the semantic net and the *conceptual affinities* between others.

It is useful to note that others are beginning to combine paradigms in a similar way. A recent attempt is described by Lim and Cherkassky [Lim & Cherkassky 1992] where semantic nets, objects and their related properties are combined in the context of more general research on knowledge representation and reasoning.

While performing this research, an important parallel observation came to light: The *language* of operations planners integrates all of these requirements equally well. The reason for this is simple: When experts communicate among each other, they are able to package concepts, however complex, in such a way as to *understand* each other. The formalization of the knowledge related to dynamic security analysis in terms of a semantic net progressively led to the realization that the resulting construction was also formally describing the *language* of dynamic security analysis.

The resulting taxonomy can therefore be interpreted as a *language*. It is useful to note that this language has semantic *and* software generality in that it is defined independently of any particular set of software tools (i.e. power flow, transient or long–term stability, short–circuit, electromagnetic transients, etc.). From an implementation perspective, the language can be viewed as a *shell* which drives a *library* of application programs. Of course, this also requires the preparation of unique bridges for each software tool to be driven by the shell. Such shells, capable of providing semantic and software generality, are also called *frameworks* [Kaplan 1992].

The remainder of this chapter gives a detailed description of the proposed taxonomy illustrated by means of the semantic net of Figs. 3–1 to 3–7.

3.3 THE SEMANTIC NET

Figures 3–1 to 3–6 present a semantic network which classifies the key dynamic security goals in operations planning where the use of algorithmic software tools such as power flow, transient or long–term stability and short–circuit is essential. Figure 3–7 addresses the types of heuristic processes that can

also be defined. There exists a very definite relationship between the algorithmic and heuristic processes and this will be explored in the discussion on heuristic processes (see 3.3.3).

Because the semantic net defines a language, one can easily build scripts, or programs, directly from the figures. Examples of how this is done are considered section 3.4.

3.3.1 GENERAL REMARKS

The most striking visual element of the semantic net resides in the choice of symbols representing the various node quantities. The triangles △ represent classes of objects. The diamonds ◇ represent specific objects; these are often used, by virtue of their relationships with classes in the net, to qualify them in some way. The rectangles □ represent properties, either of classes or objects, according to context. The ellipses ⬭ represent the language structures which relate the various classes, objects and properties of dynamic security analysis. The shaded portions represent that part of the semantic net supported by the ELISA prototype, implemented as a direct result of this work. Examples of the prototype's capabilities are considered in the next chapter.

According to a rigourous interpretation of the semantic net paradigm, the arcs themselves establish relationships. In this particular implementation, for reasons of clarity which will rapidly become apparent, the ellipses are in fact a convenient way of *tagging* the arcs which precede and/or follow them. Arcs which directly connect properties to classes simply indicate the immediate hierarchical relation between the two. To define any given dynamic security analysis process, one follows the directed arcs, retaining those language structures, classes, objects and properties required by the definition of the desired process(es), and ignoring those that are unnecessary.

A final remark. Though the formal language of classes and objects is sometimes unwieldy and the explanations may appear lengthy, the concepts are in fact quite simple and the diagrams are almost self−explanatory.

3.3.1.1 SYNTAX RULES

In general, when the **for** language element precedes a class, this means that the subsequent process(es) relate to each and every object belonging to the

class. As an example of this, the **for** which precedes the **topology** class in Fig. 3−1 implies that *all* that follows applies to every object of this class, in turn. Explaining this in more familiar terms, the subsequent process is applied to every one of a list of topologies, as is commonly seen in 3rd generation computer languages.

There are many cases where a class is not preceded by a **for**, but rather by an **of** language element: This means that *all* the objects of the class must be considered simultaneously. To illustrate this, we see that, in Fig. 3−4, if many objects of the **bus** class are specified, the acceptance criterion must be satisfied at *every* one for a case to be accepted. Classes preceded by the **under** or **with** language structures, (i.e. such as **stability** and **scale** on Figs. 3−1 and 3−5 respectively), are treated as single objects in this context: The class itself becomes the *de facto* object to be identified in the course of the process definition.

In general, only those classes preceded by a predicate (i.e. **find**, **verify** or **determine**) are unknown (see Fig. 3−1): All other objects and classes which follow must be identified in order to properly define the process.

A final point. The **zone**, **corridor** and **network** classes represent groups of buses. Hence, if a user specifies an object belonging to any one of these classes, all the buses contained in the instance are targeted. For example, if an acceptance criterion (Fig. 3−4) is applied to a **zone** object, all the buses in the targeted zone must satisfy the criterion.

3.3.2 ALGORITHMIC PROCESSES

Figure 3−1 illustrates the four key dynamic security goals which require the use of algorithmic software tools. For a predetermined list of topologies, identified as distinct objects in the **topology** class, the goals consist of **find**ing:

1) the **security** status;
2) the power transfer **limit**;
3) the **security limit** and associated worst contingency **location**;
4) the **sensitivity** of specific transfer limits to changes in network parameters.

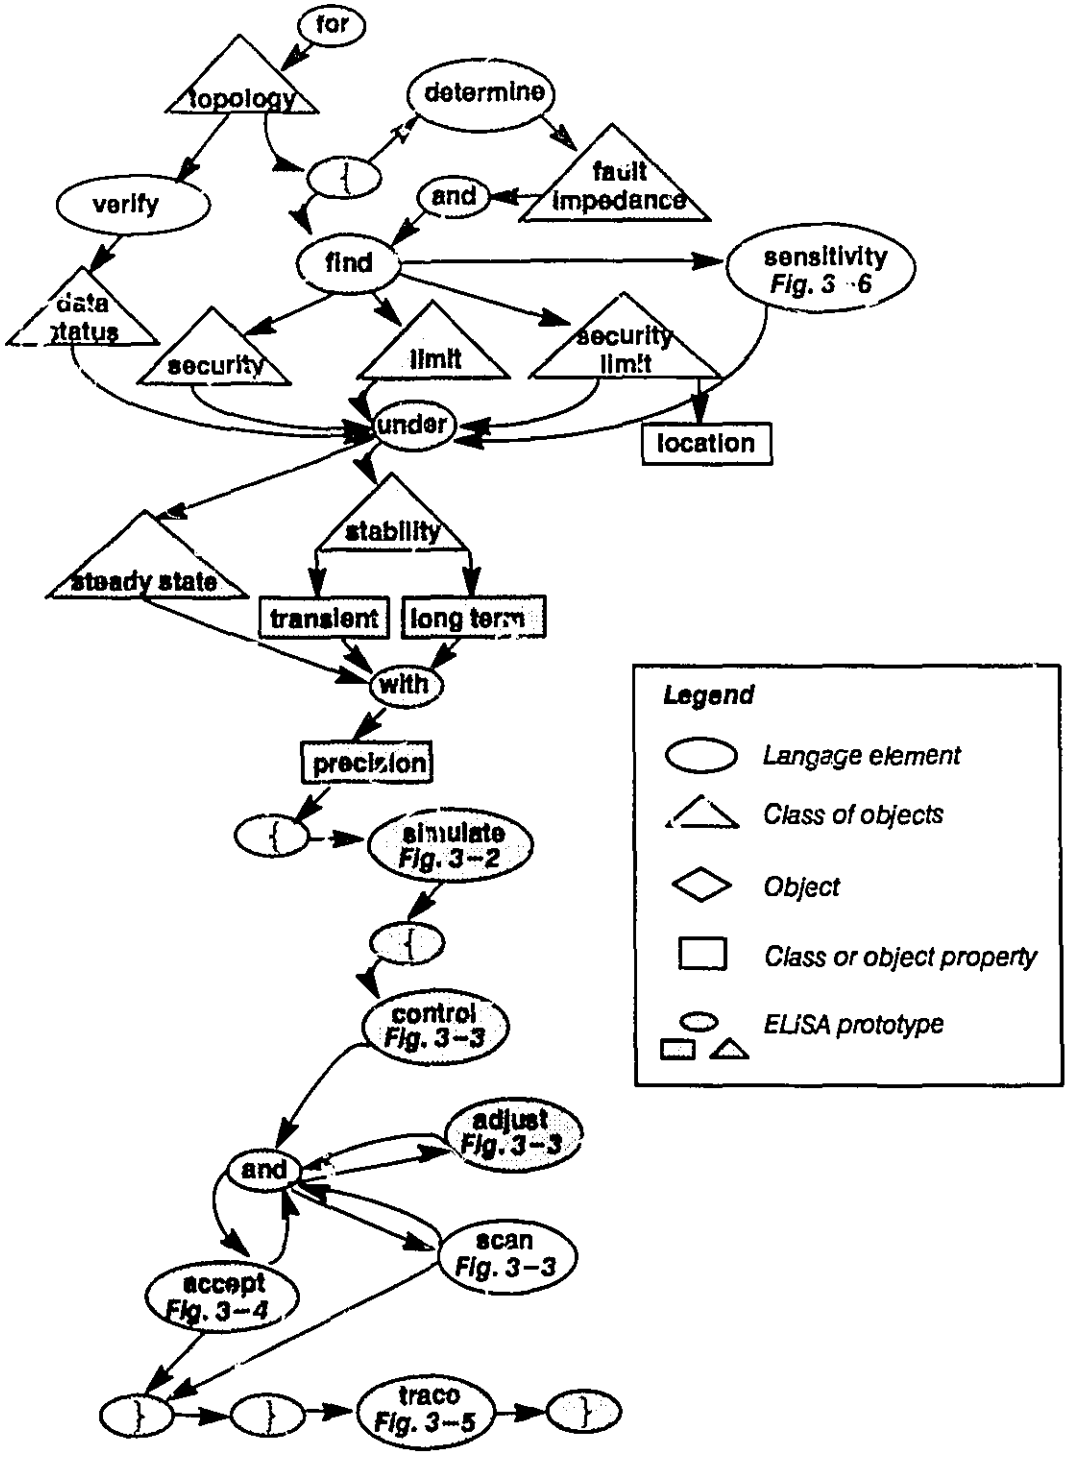

Fig. 3–1. Key diagram for algorithmic processes: The goals of dynamic security analysis.

The first goal determines whether or not a particular network topology respects acceptability criteria relative to some pre–defined contingency. The second consists of determining the power transfer **limit** with respect to some contingency. The default limit search strategy uses **controled** and **adjusted** quantities within a modified binary search. This search strategy begins by incrementing or decrementing the **controled** and **adjusted** quantities uniformly using the **initial increment** until an acceptable and unacceptable case have successively been found: This circumscribes the **limit** within the precision of the **initial increment**. Thereafter, a binary search improves this value until the desired value for **precision** (Fig. 3–1) has been found. Acceptability criteria are defined by means of the **accept** process.

The third consists of finding the **security limit**: The default strategy first determines the power transfer **limit** at different contingency locations; the most constraining power transfer limit in this set identifies the **security limit** and its **location**, the worst contingency. The object of the fourth goal is to find the **sensitivity** of power transfer **limits** as a function of a set of user–defined parameters.

As has already been pointed out, the goals of finding the security limit and worst contingency are one and the same in that determining one implies obtaining the other: Hence, the worst contingency **location** is considered a property of the security limit. In Chapter 2, an additional goal was identified: That of obtaining an appropriate power flow for a given set of injections before a limit search is undertaken. This remains a personnel–intensive activity as human expertise is still required to perform the transformation of the network and corresponding data in the form suggested in Chapter 2 (i.e. in 2.4.3 and 2.4.7). Once this transformation has been performed, an appropriate power flow can be obtained in a single step for every power transfer iteration in a stability limit search, subject to verification of the results at the end of the search.

The four goals of Fig. 3–1 are identified as classes because the result of a specific process is considered to be a unique object of the class. For example, if the process goal is to find a power transfer **limit**, the value found by the process will be placed in an appropriate object of the **limit** class, attached to a specific

object of the **topology** class, and further characterized in terms of all of the succeeding classes, objects and properties initially required to define the process (section 3.5).

In addition to the goals, one may also wish to **verify** the status of the data to be used by the algorithmic software before some high–level task is undertaken: In some cases, it may be sufficient to check the data for consistency (i.e. a coarse data filter); in other instances, data precision (i.e. a finer verification) may be an important issue. Though not a goal of dynamic security analysis by definition, this function is nevertheless a strategic corollary requirement.

3.3.2.1 OVERVIEW OF BASIC GOAL PROCESSES

To solve real problems in dynamic security analysis, one must do more than simply state the goal: The context must be defined (i.e. is the goal to be found for **transient stability**, **long term stability** or **steady state** conditions?) which in turn conditions the choice of software. If the desired goal is to **find** a **limit**, or some higher–level process, one needs to identify the **precision** with which individual limits are to be found. Additionally, the process search variables must be identified (i.e. is it a generation–load, load–load or generation–generation limit search? and where?) and the basis on which the cases are to be accepted or rejected (i.e. using voltage? frequency? applied where?). As different operations planners will likely approach problems in different ways, the semantics of the knowledge representation must be capable of accomodating a wide range of process parameters and acceptance criteria. For this reason, *semantic generality* dictates having the capability of expressing a large variety of dynamic security analysis processes, including:

1) **simulate**ing various contingencies;
2) **control**ling certain power flow inputs;
3) **scan**ing certain power flow inputs;
4) **adjust**ing other power flow inputs;
5) **accept**ing stability simulations on the basis of user–defined criteria.

As mentioned above, the semantic net also includes the definition of **steady state** security analysis processes. This refers to steady–state limit

determination by means of the well–known PV–curve limit search process [IEEE 1991, Mansour & Kundur 1991]. In operations planning, the steady–state *post–contingency* PV–curve limit search process is of considerable interest as it has been found to provide limits which correlate strongly with long–term stability limits for at least one network where voltage stability was an issue [Mailhot, Gauthier & Vanier 1993]. This explains why such processes are also included in our taxonomy of dynamic security analysis (see 3.3.2.4).

The **simulate** process, described in detail in Fig. 3–2, is used to specificy contingencies. If no explicit information is given regarding contingencies, the default contingency is simply that found in the stability data corresponding to a specific topology. When the **simulate** process is not desired, one may alternatively wish to specify the **determine fault impedance** process on the **topology** to obtain the correct SLG fault impedance as described in 2.4.8: This places the correct value of impedance in the stability data of the appropriate topology. As we shall see shortly, the **determine fault impedance** process is also a subprocess of the **simulate** process: The two are thus never specified simultaneously.

The **control** process (see Fig. 3–3 (a)) identifies the power flow quantities to be used to drive the default search strategy towards a user–defined goal. The **scan** process (also Fig. 3–3 (a)) describes a *sweep* across a range of values (i.e. using some **increment**), for example to obtain finer results after having found a **limit**.

The **adjust** process (see Fig. 3–3 (b)) is used in conjunction with **control** or **scan**: it serves to modify additional power flow inputs in order to compensate for network quantities modified by a **control** or **scan**. For example, if the **generation** of a **bus** is **control**ed, i.e. is being used to drive a search process, one may wish to **adjust** the **load** somewhere in the network in order to balance this change, as discussed in 2.4.2.

The **accept** process (see Fig. 3–4) permits the definition and application of a host of possible criteria which, when used singly or in conjunction with others at a variety of **buses**, **zones**, or **corridors**, determine whether individual network simulations are acceptable or not.

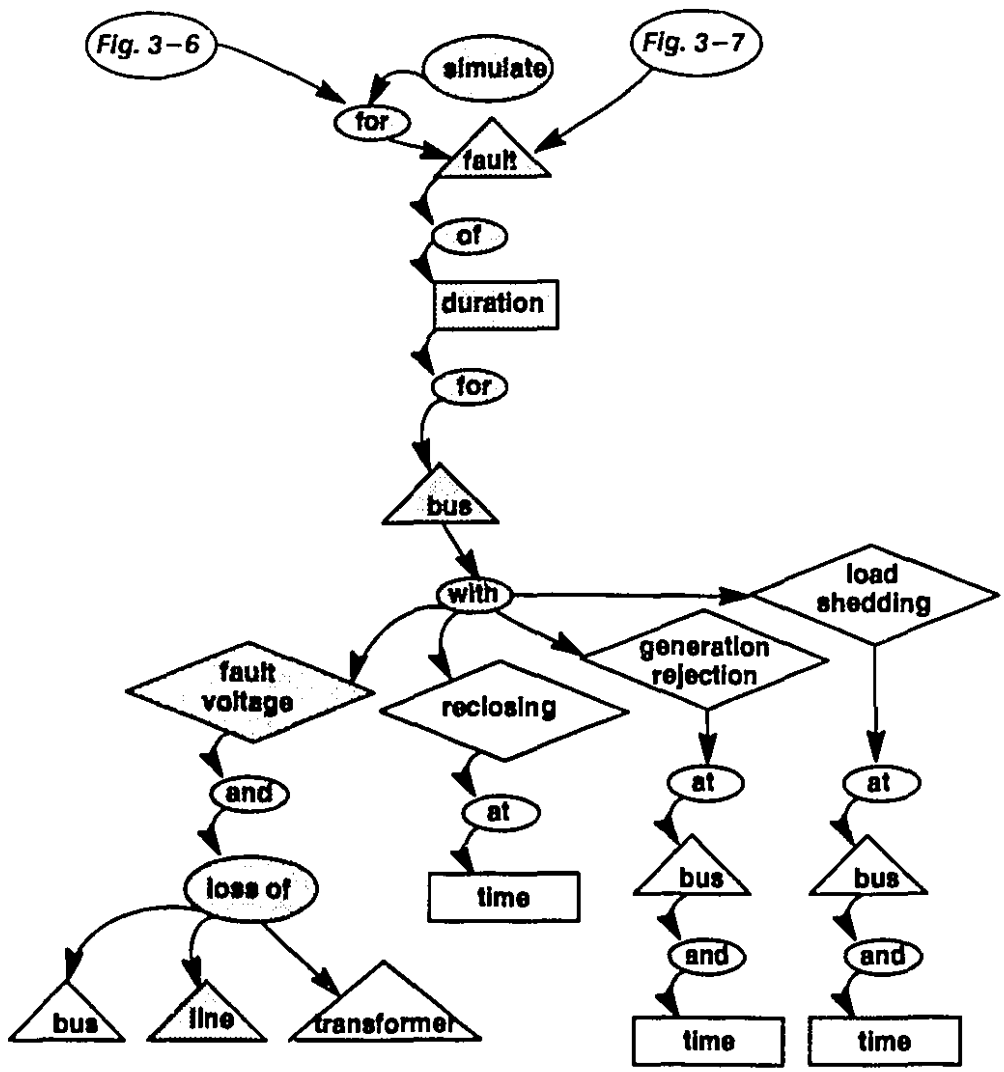

Fig. 3–2. Defining contingency characteristics: The **simulate** process.

After having specified a goal process in terms of the above, it is often desirable to obtain hardcopy or screen output of the many stability simulations performed in the course of the process. The **trace** process (see Fig. 3–5) covers this need by providing a plot of the various quantities identified within the **accept** process. This permits a visual confirmation of result interpretation.

Finally, complex sensitivity analyses can be defined using the above as elementary building blocks. These are detailed in Fig. 3–6.

As we have seen, Fig. 3–1 constitutes the key diagram for algorithmic software processes. Before considering heuristic processes, let us examine the preceding processes in detail.

3.3.2.2 THE **simulate** PROCESS: DEFINING CONTINGENCY CHARACTERISTICS

Figure 3–2 shows the many degrees of freedom that can be used to define this process. The initial **for** language structure which precedes the **fault** class implies that different fault types (i.e. three–phase, single–phase, etc.) can be identified as distinct objects of the **fault** class, each one having the property of a specific **duration**. Succeeding processes in the semantic net are thus performed for each of the **fault** class objects.

The following **for** compounds the number of contingencies to be simulated in terms of the location of the desired faults. In other words, each of the previously–specified faults is applied in turn to the list of specified locations. Consequently, if the **bus** class is specified, the faults are applied at each bus in turn. If the **zone** class is specified, each of the faults is applied, in turn, to each bus of each object of the **zone** class. The same holds for the **corridor** class.

The remaining classes, objects and properties further qualify the contingency scenario in terms of topology changes and more complex events. For example, the desired **fault voltage** (as outlined in 2.4.8) must be specified at each of the targeted buses for SLG faults (the SLG fault impedance is calculated as described in 2.4.8), and equals zero for a three–phase fault. If must be stated that the fault voltage approach may not necessarily be appropriate for other types of faults, and a rigourous approach involving the use of short–circuit software may be necessary in such instances. One may also wish to simulate the

simultaneous loss of one or many **buses**, **lines** or **tranformers**, as is normally the case when clearing a faulted element.

If **reclosing** on a fault is to be simulated, the **time** at which this reclosing is to take place must be specified. The same holds for **generation rejection** or **load shedding** strategies where the list of targeted buses must be identified, including the **time** at which the action is to take place.

One final point. Network actions such as **reclosing**, **generation rejection** and **load shedding** qualify specific fault scenarios. They are considered objects rather than **fault** properties in that *additional* attributes must be specified in each case, such as the **time** at which they may take effect, and the **buses**, **lines** or **transformers** which may be involved.

3.3.2.3 THE **control**, **scan** and **adjust** PROCESSES: DRIVING THE POWER FLOW

Figure 3–3 illustrates the degrees of freedom required to express the many search strategies which depend on modifying power flow operating points.

When the goal is to **find** a **limit** or some higher level process, one must **control** either **generation** or **load**. In more complex networks, one may wish to have direct **control** on the **transfer** over a **line** or a **corridor**: This can be the case for HVDC links [Kimbark 1971] or FACTS [Hingorani 1993] power control technologies. In either case, **control** can be exercised on a set of **zone** or **bus** objects, or on the entire **network**.

When many objects are designated for **generation** or **load** changes in either the **zone** or **bus** classes, each one must be given a relative **weight** in relation to the others: One cannot presume that the generation or load is to be divided equally among the different objects. However, in a single **zone**, the distribution is always performed on a *pro–rata* basis. The same holds for objects attached to the **line** and **corridor** classes. It must be stressed that, conceptually, the semantic network permits sets of **buses**, **zones**, **lines** and **corridors** to be **controled** either separately, or any in desired combination.

When the goal is to **find** a **limit** or a **security limit**, the **initial** search step

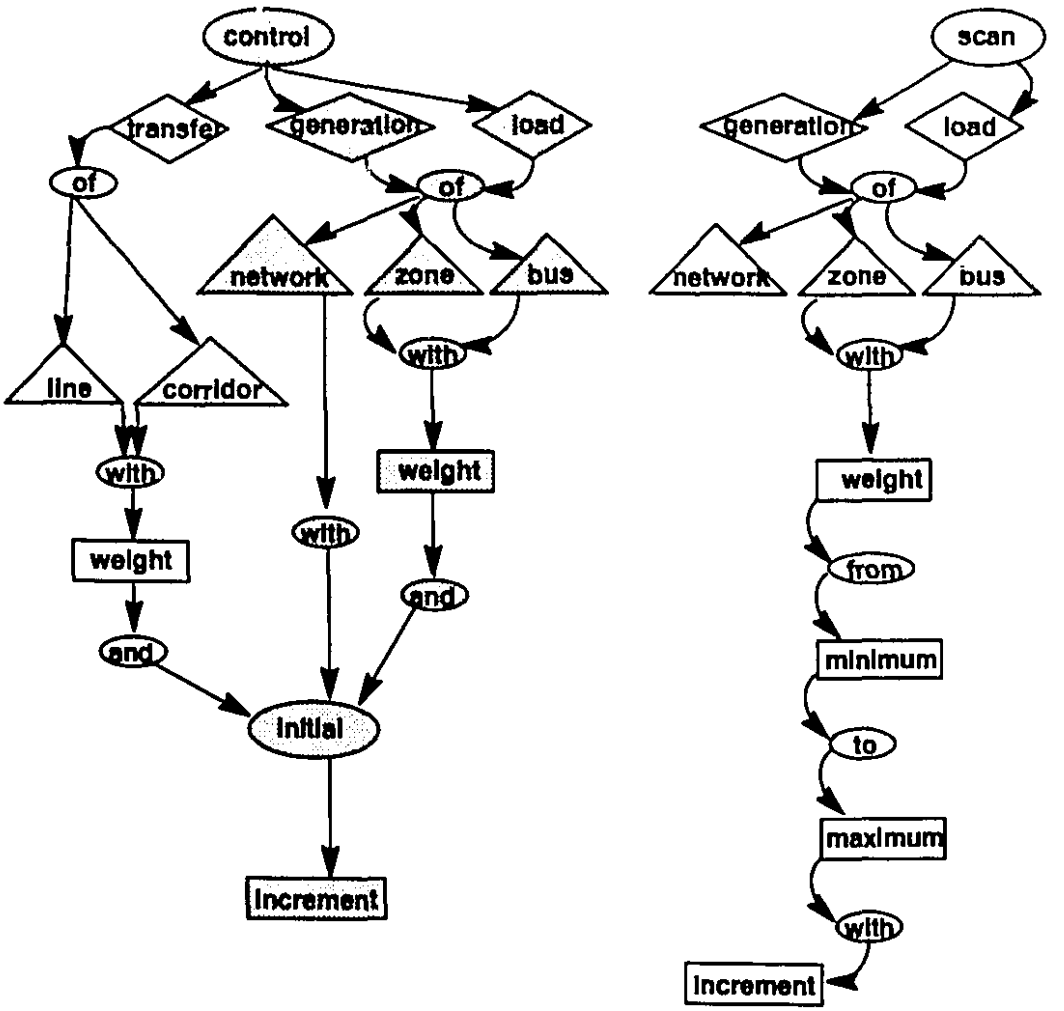

Fig. 3–3 (a). Driving power flow inputs:
The **control** and **scan** processes.

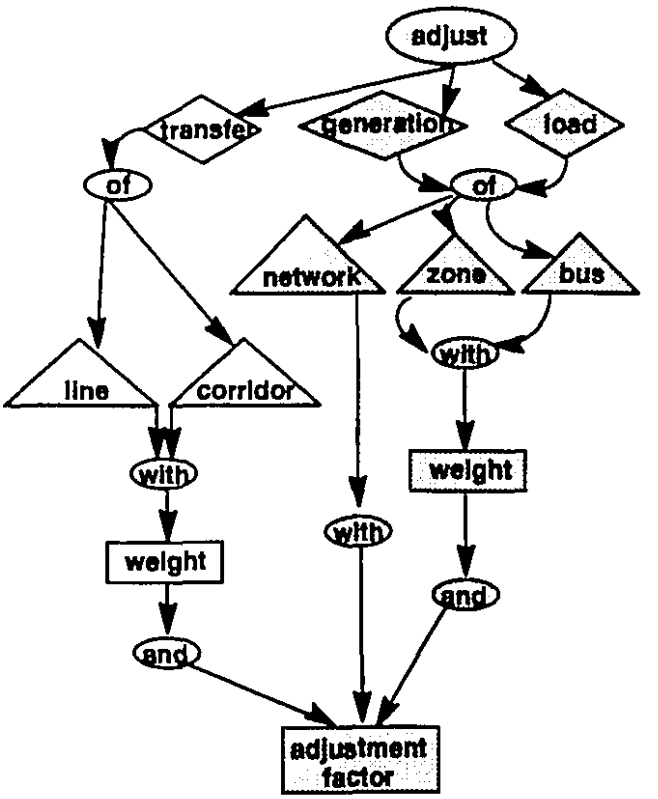

Fig. 3–3 (b). Driving power flow inputs:
The **adjust** process.

or **increment** must be specified (in MW) for the default strategies that are employed by these goals. If no increment is specified, a default **increment** is provided. The **control** process always reacts in the appropriate fashion to the **accept** process analysis: If a case is acceptable, it increments; if not, it decrements. The increment or decrement value itself is provided by the default search strategy. When no higher–level search strategy is specified because the **control** is not used in the context of a **find**, it uses the **initial increment** (see section 3.6).

Having identified the control quantities, one must also identify the corresponding **adjust** quantities. Everything previously said for **control** also holds for the **adjust** command.

The **scan** command describes a *sweeping* process, performing simulations at regular, predefined **increments** rather than attempting a search: Consequently, a **scan** has no need for an **accept** criterion. There are two different types of **scans**. First, when a **scan** is requested immediately following a limit search, the **minimum** and **maximum** represent values relative to the limit. If no minimum and maximum values are specified in this case, the default is simply to **scan** between the closest acceptable and unacceptable cases. Second, if the **scan** is asked for on a stand–alone basis (i.e. not preceded by a search), the **minimum** and **maximum** must be considered relative to the base (i.e. the starting) case. In all instances, an appropriate **increment** must be specified. Of course, when performing a **scan**, some quantity(ies) must also be appropriately **adjusted**.

A final point. The **generation**, **load** and **transfer** objects serve to characterize individual buses or lines. They are considered objects rather than **bus**, **zone** or **network** properties in that *additional* attributes must be specified, such as **weight**, **minimum**, **maximum** or **increment** for the context to be properly identified.

3.3.2.4 ACCEPTANCE CRITERIA: THE **accept** PROCESS

The **accept** command describes a wide variety of rules which are the foundation of higher–level processes. As outlined in 2.4.9, the basic approach to evaluating a simulation's acceptability is as follows: certain network quanti-

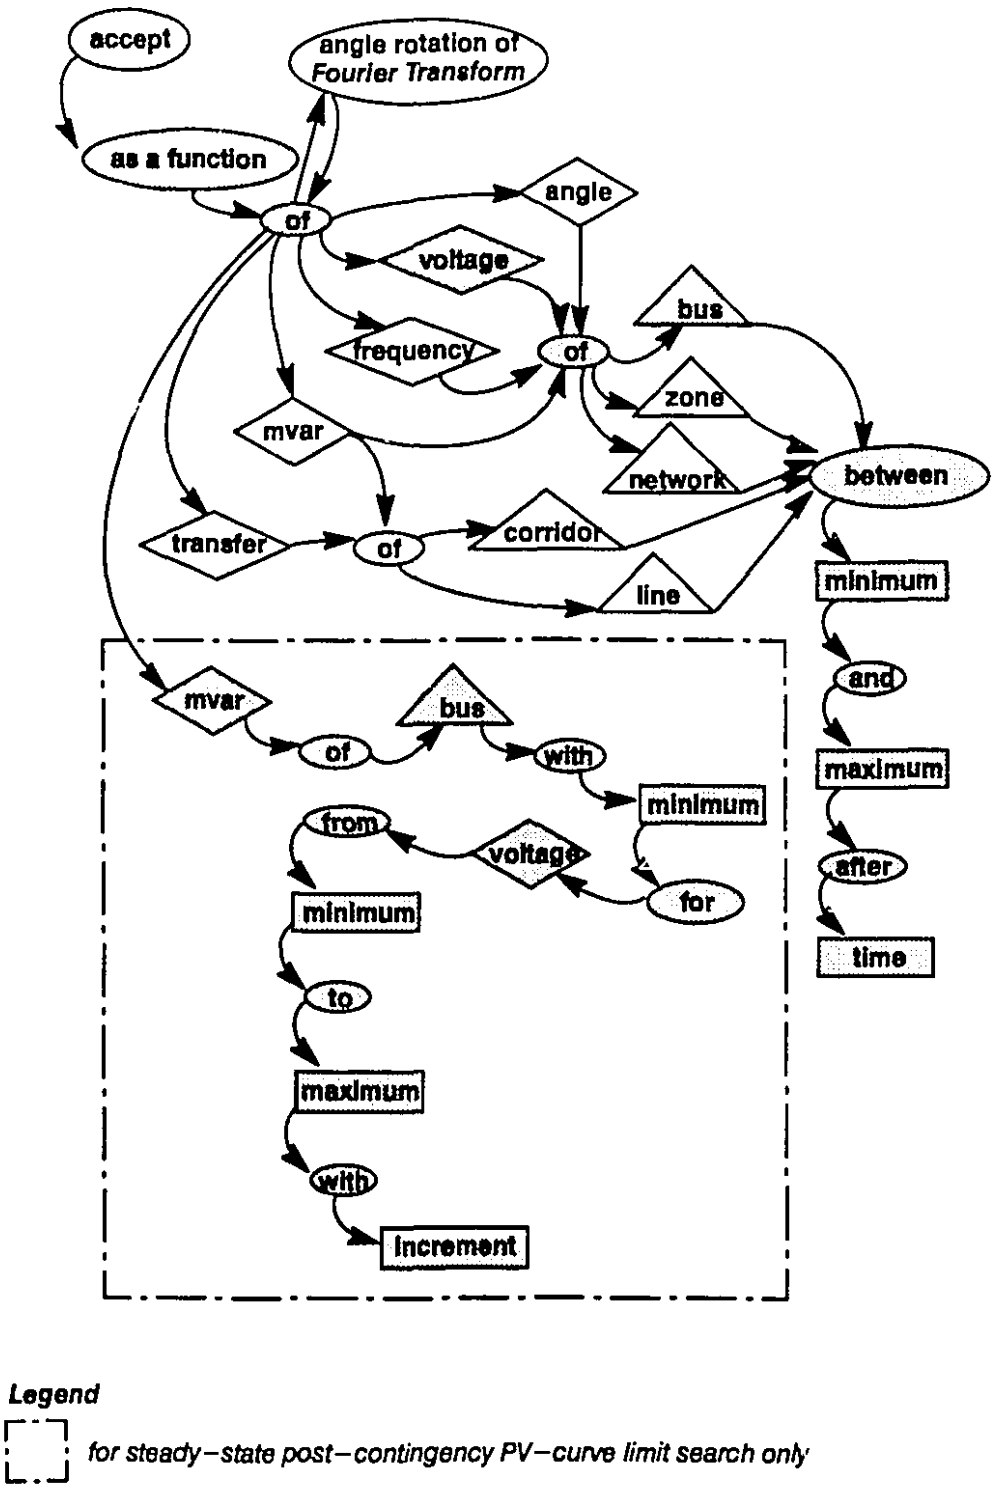

Fig. 3–4. Acceptance criteria:
The **accept** process.

ties, chosen by an operations planner, must remain within a range of acceptable values, defined at the planner's discretion and based primarily on experience.

Consequently, a transient or long–term simulation is **accepted as a function of** either the **voltage** magnitude, voltage **angle**, **frequency** or **mvar** output/input at a set of objects of the **bus**, **zone** or **network** classes. As has already been pointed out, the **of** which precedes the classes implies that when multiple objects are attached to a class, they are taken as a set, meaning that the criterion at each and every one must be passed for the case to be acceptable. Criteria can also include the **mvar** or power **transfer** along **corridor**s or **line**s, particularly tie lines. In all cases, a **minimum** and a **maximum** is specified which must be respected **after** a certain **time**. Chapter 6 of this thesis proposes a stability criterion in terms of the rotational angle behaviour of the Fourier transform of some time–varying network quantity, such as voltage or frequency. This can also be accomodated by the semantic net and is consequently shown in Fig. 3–4.

There is no conceptual constraint on the mixing of criteria for **transient** or **long term** stability. One may wish to formulate a complex criterion where the **voltage** at several **bus**es, and at every bus within a **zone**, is to be monitored in addition to the **frequency** at another set of **bus**es and the power **transfer** on a particular tie **line**. However, the semantic net also defines the acceptance parameters of the **steady state** PV–curve limit search process [IEEE 1991] and these cannot be combined with **transient** or **long term** stability criteria except by means of special filters. Such filters are considered in 3.6.3.

The steady–state post–contingency PV–curve limit search criterion described in Fig. 3–4 requires identifying the **bus** at which the reactive power Q generated for different values of **voltage** V is to be compared against a **minimum** value of **mvar**: This is repeated for every value of power P generated by the **controle**d location which has been selected to drive the limit search. The **voltage** V is varied between user–defined **minimum** and **maximum** values in specified **increments**, and the reactive power Q (i.e. of the bus where voltage is varied) is compared with respect to a **minimum** admissible value of **mvar**s, also user–defined. Rigourously speaking, **minimum mvar** value should be zero; however, one may which to use a more conservative value to provide margin, which therefore requires specifying a value less than zero. The PV–curve

limit determination criterion is therefore as follows: If Q drops below the **minimum mvar** at some value of voltage V, P is less than the limit; if Q is greater than the **minimum mvar** for all values of V, P is greater than the limit. Employing this criterion within a binary search strategy, the PV−curve limit can easily be found (see section 4.5 for greater detail).

A final point. The **angle, voltage, frequency, mvar** and **tranfer** objects are physical quantities attached to individual buses or lines. Here again, they are considered objects rather than **bus, zone, network, line** or **corridor** properties in that *additional* attributes must be specified, such as the **minimum** and **maximum** acceptable values, and after what **time** the criterion is to be applied.

3.3.2.5 HARDCOPY OR SCREEN OUTPUT: THE **trace** PROCESS

In many cases, it is important to obtain some trace of **transient** or **long−term stability** simulation results in order to verify the analysis and to acquire deep knowledge about various network phenomena. Fig. 3−5 shows the semantic net required to define the **scale** of the various quantities.

The basic premise here is that the quantities which define the acceptability criterion should be output for the purpose of verification. Consequently, one need only identify the desired scale for the plotting of quantities, not the quantities themselves.

3.3.2.6 SENSITIVITY ANALYSIS

Sensitivity analysis is a large area which can only be circumscribed in a general way. Figure 3−6 illustrates the basic concept underlying sensitivity analysis: The main objective is to observe how **limit** values vary in terms of certain network parameters, such as the number of **SVC**s and their location, the **generation** or **transfer** at some related point in the network, the transient **load characteristic** of the network, system **voltage** profile, any number of different **control** system **settings, fault duration** and **fault location**. Other network quantities may also be of interest. The final result is a table which shows how the limit varies in terms of the chosen parameters, stored as a specific object of the sensitivity class.

A basic premise in this figure is that **control**ed and **adjust**ed quantities are identified downstream in the semantic net, in the same way as one would nor-

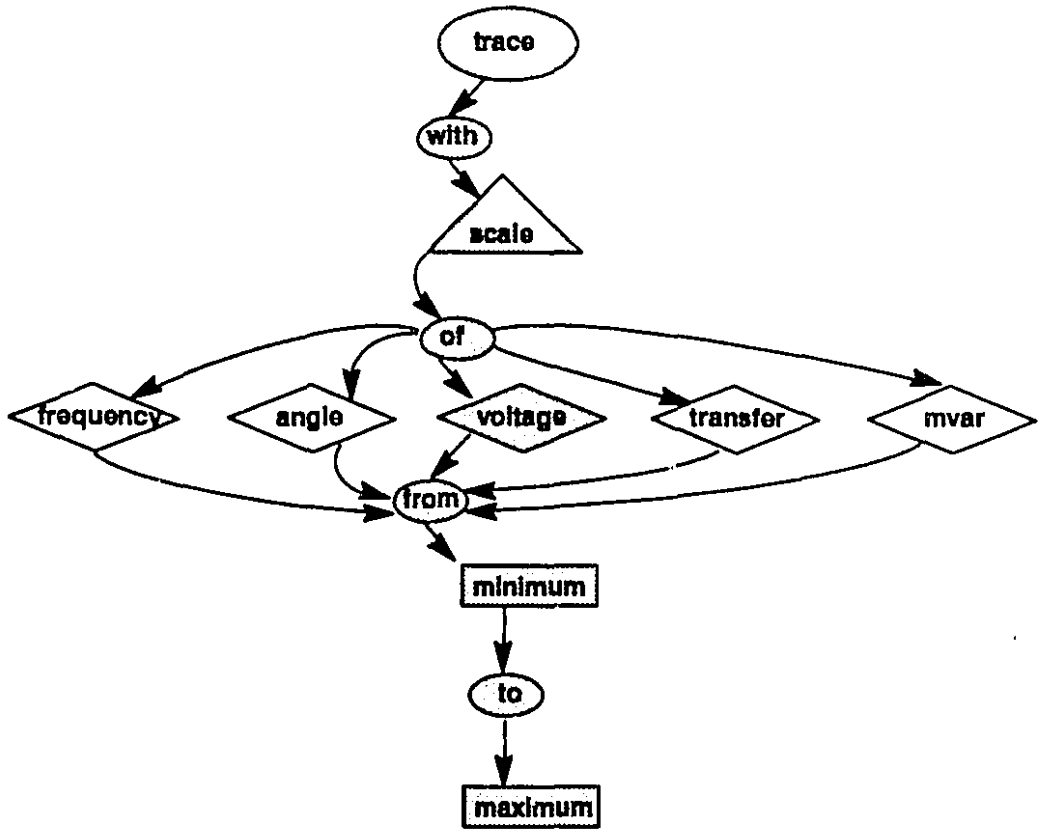

Fig. 3–5. Plotting of simulation results:
The **trace** process.

mally define them for an ordinary limit search. Consequently, only those network *parameters* to which the sensitivity is to be determined are identified.

As before, **FACTS, SVC, generation, transfer, voltage, load characteristic, control setting, fault duration** and **fault location** are physical quantities attached to different individual network components. They are again considered as objects rather than **bus, zone, network, line** or **corridor** *properties* in that *additional* attributes must be specified for them to be useful within the context of the process.

When considering sensitivity analysis, one must identify the range of values between which the parameters are to be varied: Consequently, such quantities such as the **minimum, maximum** and **increment** must be specified. The loca-

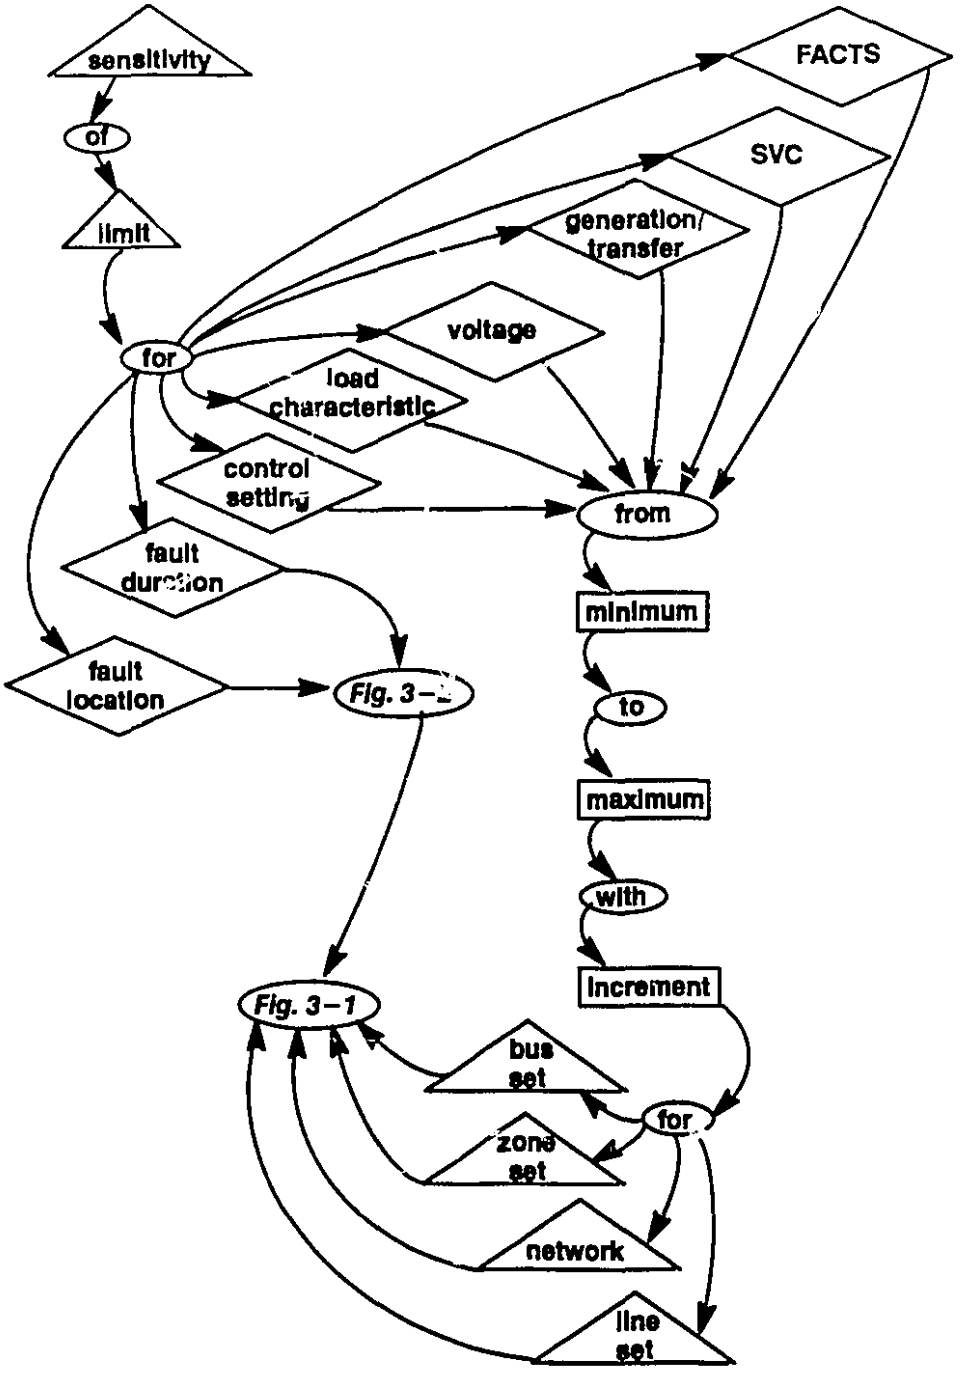

Fig. 3–6. Characterization of typical sensitivity analyses: The **sensitivity** process.

tions at which the parameters are to be applied must also be identified, and these are specified by means of the **for** which precedes the **bus set**, **zone set**, **line set** and **network** classes, causing the process to be repeated for each object in turn.

The **bus set**, **zone set** and **line set** classes are similar to the **bus**, **zone** and **line** classes previously encountered, though they also differ in important ways. An object of the **bus set** class is a *group* of buses which are to be considered *simultaneously* as a set; the same holds for **zone set** and **line set**. When an object of the **bus set**, **zone set** or **line set** classes includes a single element (i.e. a single bus, zone or line respectively), then they are indistinguishable from objects of the **bus**, **zone** or **line** classes.

There are two usages, depending on the type of parameter, that is, either continuously variable or discrete parameters. For example, when performing sensitivity analysis on **voltage**, **load characteristic** or some **control setting**, all of which are continuous quantities, an object of the **bus set** class will identify a group of buses where the parameter must be changed *simultaneously*. This also holds for **generation** sensitivity analysis, except that whenever **generation** is incremented at some specific **bus set** or **zone set**, an equivalent **adjust** process must also take place: This is taken to be the **adjust** process explicitly defined downstream.

The second usage concerns sensitivity analysis on discrete elements such as **SVC**s or **FACTS** devices (i.e. variable series compensation, dynamic phase angle regulators, etc.). Such devices are not found at every bus or line, and device distribution, for a fixed number, is an important issue. The **bus set**, **zone set** or **line set** classes permits this dimension to be taken into account through a unique feature which generates a *set of combinations of discrete parameters* in direct relation with the total **incremented** value. Let us explain by considering an example.

Suppose that one wishes to perform a sensitivity analysis of the impact of up to 2 additional **SVC**s distributed among two buses, say 780 (i.e. Nemiscau) and 782 (i.e. Albanel) on the Hydro–Québec network. We first identify an object of the **bus set** class as being *780, 782*. We then set the **minimum** to 0

TABLE 3–1

Example of the Use of the **bus set** Class in SVC Sensitivity Analysis

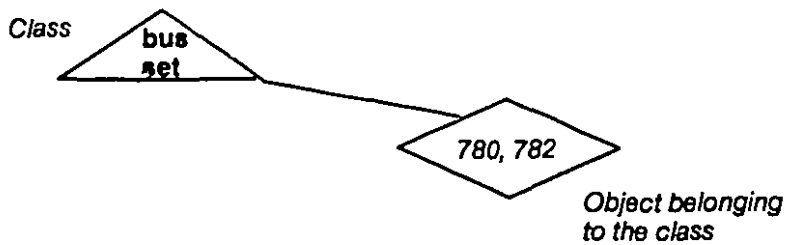

Case	*SVC total*	*SVCs at bus 780*	*SVCs at bus 782*
1	0	0	0
2	1	1	0
3	1	0	1
4	2	2	0
5	2	1	1
6	2	0	2

and the **maximum** to 2, **incrementing** one (1) SVC at a time. The **bus set** class therefore generates, for each incremented *total* of SVCs, all the possible combinations distributed among the two buses, as shown in Table 3–1. Though only SVCs are mentioned here, the same holds for SCs and shunt reactors. For active series power control concepts represented by the emerging **FACTS** technologies [Hingorani 1993], the same kinds of processes hold using the **line set** class.

Sensitivity analyses concerning **fault duration** and **fault location** are far more straightforward and their details have already been covered in the discussion of Fig. 3–2.

3.3.3 Heuristic Processes

A number of heuristic processes can be defined for the operations planning environment. However, the fundamental premise here is that there exists a large database of cases on which heuristic methods can be applied in order to answer the types of questions that are of interest to dynamic security analysis. On average, even simple heuristics can provide good results when the database is large. However, the smaller the database, the lower the expected precision.

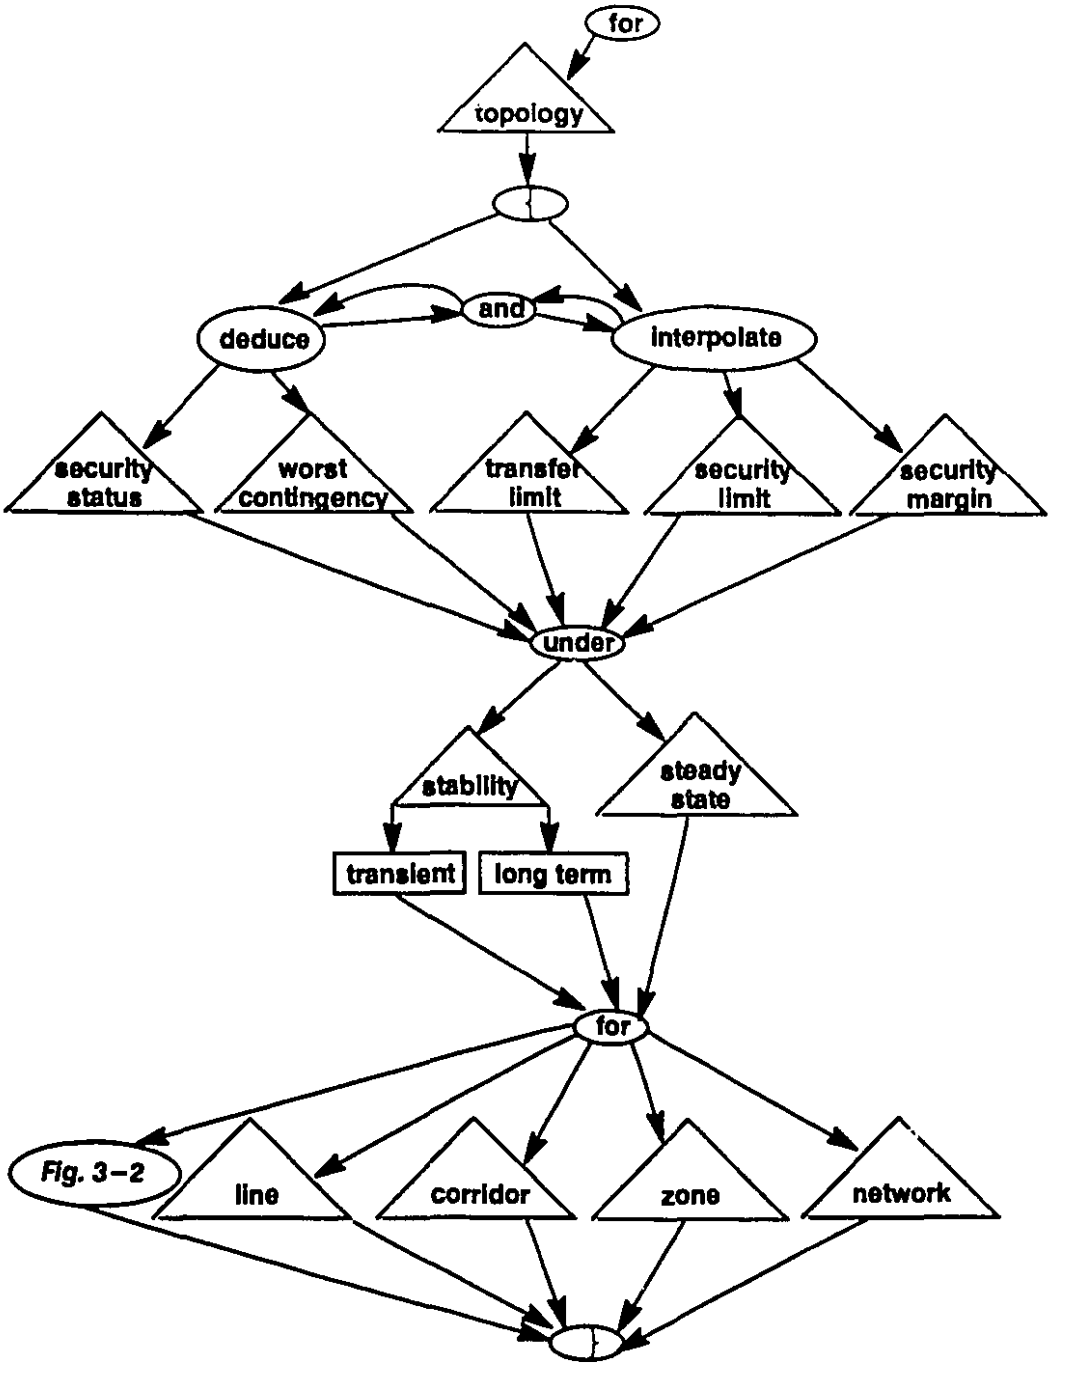

Fig. 3–7. Key diagram: Heuristic processes.

Of course, here we probe the direct cause–and–effect relation between algorithmic and heuristic processes. Heuristic processes are dependent on the capacity of algorithmic processes to generate a sufficiently large number of stability simulations, transfer limits, worst contingencies and security limits, these in turn reflecting primarily those of pre–defined umbrella line topologies (i.e. Nd_i networks, see 2.4.4) and various voltage support configurations (i.e. Nd_{ij} networks). If the data base is very large, the query process might even be capable of performing heuristic sensitivity analysis, provided that an appropriately detailed classification of the security data has been achieved.

An additional issue must be addressed. In the light of present trends in EMS technology which tend towards on–line security analysis, preferably using real–time simulations, is there a future for heuristic methods? Though such questions are difficult to answer with any degree of certainty, I believe that the answer is a qualified "yes".

In the short term, if heuristic methods can be developed and mechanized quickly enough, the manual compilation of transmission limit tables or decision trees from off–line results can be replaced by means of mechanized heuristic methods accessing a database of off–line simulation data. This would eliminate the tedious human processing and transcription of dynamic security analysis results [Avramovic & Fink 1991]. These results could then be verified and forwarded to the system operator.

In the medium term, when such an approach will have been borne out, its natural extension is to integrate the same heuristic methods directly to EMS technology, and on–line verification of system security could be undertaken at regular intervals from a database of off–line simulation data.

In the long term, even if real–time stability simulations become an economic reality, such techniques might remain extremely useful in estimating an appropriate starting case or an initial increment from which a limit search can be undertaken, or by circumscribing the set of locations at which one might expect the worst contingency in security limit estimation, thereby reducing the time required to perform such functions in real–time. However, the greatest potential of heuristic frameworks may be to provide a foundation for the devel-

opment of advanced decision support software in the on-line environment, capable of identifying preventive, corrective or restorative control measures.

3.3.3.1 PRELIMINARY CONSIDERATIONS: HIGH-LEVEL DATABASE STRUCTURE

The heuristic processes identified in Fig. 3-7 are based on the hypothesis that **security status, transfer limits, security limits** and **worst contingency** information are stored, at the highest level, in terms of:

1) network line topology (Nd_i networks);
2) active voltage support component distribution (Nd_{ij} networks); and
3) contingency C_k.

This is according to the general approach outlined in 1.6.2.1.

A final point. The basic heuristic process is a query. Given a specific network topology, the query is identified, such as **deduce**ing **security status** (is the indicated topology secure?) or the **worst contingency**, or **interpolate**ing the **security limit**, the **security margin** or the **transfer limit** for a contingency at a specified location. Such queries can be made under the assumption that the data reflects **transient** or **long-term stability** simulation data, or **steady-state** post-contingency data.

3.3.3.2 **Interpolate**ING **security limits**

for a list of **topologys**, one may wish to **interpolate** a **security limit**, for a list of **lines** or **corridors**, or for every line in a **zone**. When the **security limit** is to be determined for a **line** which is itself part of a specific corridor, one determines first the corridor security limit, then the corresponding value for the line.

The general approach is as follows. A closest-fit search is performed on all umbrella line topologies: This yields the one which subsumes the target topology and restricts the primary search space in terms of Nd_i networks. A number of strategies can be used to accomplish this, many of which are outlined in [Stott, Alsaç & Monticelli 1987].

The number of active voltage support elements (SVCs, SCs, etc.) in the targeted corridor are then identified: This circumscribes the search space in terms

of a family of Nd_{ij} networks. Small variations in the number of voltage control elements may be tolerated in adjacent corridors provided that they are few in number and that the network is subsumed by the database data, or that it is clear that they have no impact on the limit. Even large variations may be tolerated in distant corridors provided that it is clear that their impact is minimal.

The next step is to consider the actual *distribution* of voltage support elements. If an identical distribution is found, the security limit is extracted with no further ado. If it is not, the closest configurations are once again identified and a conservative value for **security limit** is estimated by interpreting trends in the data.

3.3.3.3 **deduce**ING **worst contingency**

deducing the **worst contingency** requires a process identical to the above except that, having found the **security limit**, the associated location consequently represents the **worst contingency**.

3.3.3.4 **deduce**ING **security status**

deducing security status means that one is interested in knowing whether or not the targeted topology and its particular operating point are contained within a secure region of operation with respect to applicable security limits. The focus of such a request can be the power transfer a) over a specific **line** or **corridor**, b) for all the lines in a **zone** or c) on all of the lines of the entire **network**.

To perform a comparison of operating point and security limit for a single object, such as a line, the first step is to determine the **security limit** of this object as previously outlined. Then, a simple comparison establishes **security status**: If the value of the security limit is greater than or equal to the topology operating point, the topology is secure.

When the **security status** of multiple elements is requested (i.e. many **lines**, the global value for many **corridors**, or for every line in a **zone**) each individual object is considered in turn. When considering a class such as **corridor** or **zone** where each consist of many lines, each line is considered in turn within each corridor or zone.

3.3.3.5 **interpolate**ING **security margin**

For **interpolate**ing **security margins**, the process is identical to that of **interpolate**ing **security limits** except that, having found the **security limit**, the operating point of the targeted topology is obtained and subtracted from the limit value: This gives the desired margin.

3.3.3.6 **interpolate**ING THE **transfer limit**

When one wishes to **interpolate** a specific **transfer limit**, contingency and location must be identified using the semantic net of Fig. 3–2, though very few fault *types* (i.e. SLG, three–phase, etc.) are considered in practice. From this point, the search progresses in a manner identical to that of the **security limit** search except that, having identified the closest Nd_{ij} topologies, the corresponding transfer limit data is extracted. Data trends must then be analyzed in terms of contingency type, location and topology–dependent transfer limits. Provided that sufficient simulation data is present, one can estimate conservative values.

3.4 USING THE SEMANTIC NET TO DESCRIBE DYNAMIC SECURITY ANALYSIS PROCESSES

We have seen that the semantic net provides a formal structure for the description of dynamic security analysis processes. In the present section, two examples are presented by means of the proposed semantics and syntax: This exercise plainly illustrates the relationship between the formal description of dynamic security analysis processes, presented in the form of a semantic net, and the language of operations planners.

To build process descriptions, one combines successive elements found on the semantic net. The two examples illustrate all of the basic ideas: As a path is followed down the net, individual elements are used as required, or simply ignored. Once again, it must be stressed that the formal language of classes and objects may sometimes appear unwieldy and the explanations lengthy, but the concepts are simple and the scripts themselves are essentially self–explanatory.

One small point. For the sake of clarity in the presentation, some syntax elements have been abbreviated. For the same reason, certain conjunctions (i.e. **and, with**, etc.) which are obvious in the context have been left out altogether.

```
for   topology
      hq1
      hq2
  {find limit under  transient stability with   precision
                     st600                      100
        {control generation of    bus   weight with   init_inc
                                  49    .50            400
                                  50    .25
                                  64    .25
        and adjust load of network with  adjust_factor
                                                 0.9
        and accept as a function of voltage of bus  min  max  after
                                                 702  0.96 1.0  400
                                                 703  0.95 1.0  450
                                                 714  0.94 1.0  350
        }
  and trace with scale of voltage   min  max
                                    0.9  1.1
  }
end
```

Fig. 3–8. Description of a typical generation–load limit search process by means of the semantic net.

3.4.1 A TYPICAL GENERATION–LOAD TRANSFER LIMIT SEARCH

Figure 3–8 shows a simple script describing a typical generation–load transfer limit search. This script describes a limit search for two topologies. The generation at a group of buses drives the search process, the load is adjusted across the entire network and, for a case to be acceptable, the voltage at buses 702, 703 and 714 must satisfy the criterion specified at each bus. What follows now is a more detailed, formal explanation of the script in terms of the classes and objects of the semantic net.

```
for topology
    hq1
    hq2
```

identifies the objects of the **topology** class which are to be treated in turn. Here we see two topologies, hq1 and hq2. Typically, each object name points to two different types of data, power flow and transient stability data, the latter including the contingency to be applied.

```
{find limit under transient stability with precision
                    st600                    100
```

uses elements of Fig. 3–1 to identify a goal (i.e. the power transfer limit) and to specify that this goal applies to transient stability simulations using the st600 transient/long–term stability software. As mentioned above, the goal is found in turn for each topology in the list because a **for** precedes **topology**. The **precision** of the desired **limit** is specified as 100 MW. The opening of the parenthesis { indicates that everything contained within applies, in turn, to each and every topology. As we have also seen, the default search strategy is a modified binary search. The result of each individual simulation (i.e. is it acceptable or not?) is stored in a different object of the **stability** class.

```
{control generation of bus  weight with  init_inc
                    49   .50          400
                    50   .25
                    64   .25
```

uses elements of Fig. 3–3 to show how the **generation** object is used to control generation of **bus** class objects 49, 50 and 64 to drive the process towards the goal. A specific **generation** object is attached to each of the targeted **bus** objects: The **weight** property belongs to each **generation** object, so that each specific value is stored there. The opening of the second parenthesis { indicates that everything contained within describes the processes required by the **find limit** goal. The **limit** search begins with an **initial increment** of 400 MW, switching to a binary search when the search space has clearly been delimited.

and adjust load of network with adjust_factor
.9

As we have seen, a **control** command must always be paired to an **adjust** command so that changes, for example, of generation, are balanced by equivalent load or generation changes elsewhere. Also, as no **simulate** command is provided in the script, the default contingency is that found in the topology stability input data. In this case, **load** is **adjust**ed, and the **adjust_factor** property, which belongs to the **load** object, is given the value of 0.9.

and accept as_a_function of voltage of bus min max after
702 0.96 1.0 400
703 0.95 1.0 450
714 0.94 1.0 350

}

uses elements of Fig. 3–4 to build an acceptability criterion. In this particular case, we see the acceptable voltage range for each of the specified buses. Formally speaking, the acceptable **min**imum and **max**imum voltages at each bus are stored in appropriate properties of a **voltage** object belonging to every one. If the **zone** class had been specified rather than **bus**, every bus in the specified zone would have had to satisfy the criterion. The closing parenthesis indicates that the **find limit** goal process is described by the preceding.

and trace with scale of voltage min max
0.9 1.1

}

```
for   topology
        top1
        top2
        top3
   {find limit under  transient stability with   precision
                        st600                      100
      {simulate for fault of  duration
                                4
                                6
                                8
         for   bus  with   fault_volt  and   loss_of_line
               720            .625          720 780 1
               723            .707          723 782 1
               780            .655          780 713 1
               782            .661          782 783 1
               713            .649          713 714 1
               783            .653          783 783 1
               731            .739          731 717 1
               714            .658          714 770 1
               717            .677          717 702 1
            {control generation of    zone weight with   init_inc
                                       59        1          800
            and adjust load of network with  adjust_factor
                                                  0.9
            and accept as a function of voltage of bus  min  max  after
                                                  702   0.95  1.05   500
            }
      }
      and trace with scale of voltage   min   max
                                        0.9   1.1
   }
end
```

Fig. 3–9. Building large tasks by means of nested **for** loops.

provides instructions regarding the scale of the hardcopy or screen plot which is desired (Fig. 3–5) so that graphical output of the results at every bus are plotted as shown in Fig. 2–2. The correct scale is important so that observable phenomena can be properly monitored and verified. The closing parenthesis indicates that this completes the desired process for a given topology.

3.4.2 COMPOUNDING COMPLEXITY: NESTED **for** LOOPS

The main feature of the next example, illustrated in Fig. 3–9, is that of multiple **for** loops nested within one another. In this example, the goal is to find a transfer **limit**, and the limit search process is essentially described in terms of driving the **generation** in **zone** 59, adjusting **load** on the whole **network** with an **adjust_factor** of 0.9, and monitoring **voltage** at a single **load bus** (i.e. bus 702) to ensure that it remains within emergency limits after 500 cycles. Conceptually, much of this resembles the preceding example. The main difference relates to the contingencies to which the basic search process will apply.

This script shows three different **for** language elements. The lowest level **for** (i.e. **for bus** ...) relates to fault location and fault voltage. In this case, we wish to apply a SLG fault with subsequent loss of line, in turn, at every one of the nine buses found on the James Bay corridor: This will enable us to determine the worst contingency and the corresponding security limit. The *preceding* **for** (i.e. **for fault**...) specifies that this be done for three different fault durations. The *first* **for** (i.e. **for topology** ...) further specifies that this be executed, in turn, on three different topologies. All in all, this short script describes the execution of 9 x 3 x 3 = 81 distinct transfer limit searches.

A final remark. Individual power transmission lines are often identified by the buses at their extremeties. Because a number of lines may be in parallel between two buses, a *circuit number* is also needed to complete the line identification. For example, in Fig. 3–9, the line to be lost for a fault at bus 720 is the line between buses 720 and 780, *circuit number* 1.

3.5 CLASS–OBJECT NETWORKS

When examining the semantic net, the question naturally arises as to whether or not it can be viewed as defining hierarchical relationships between

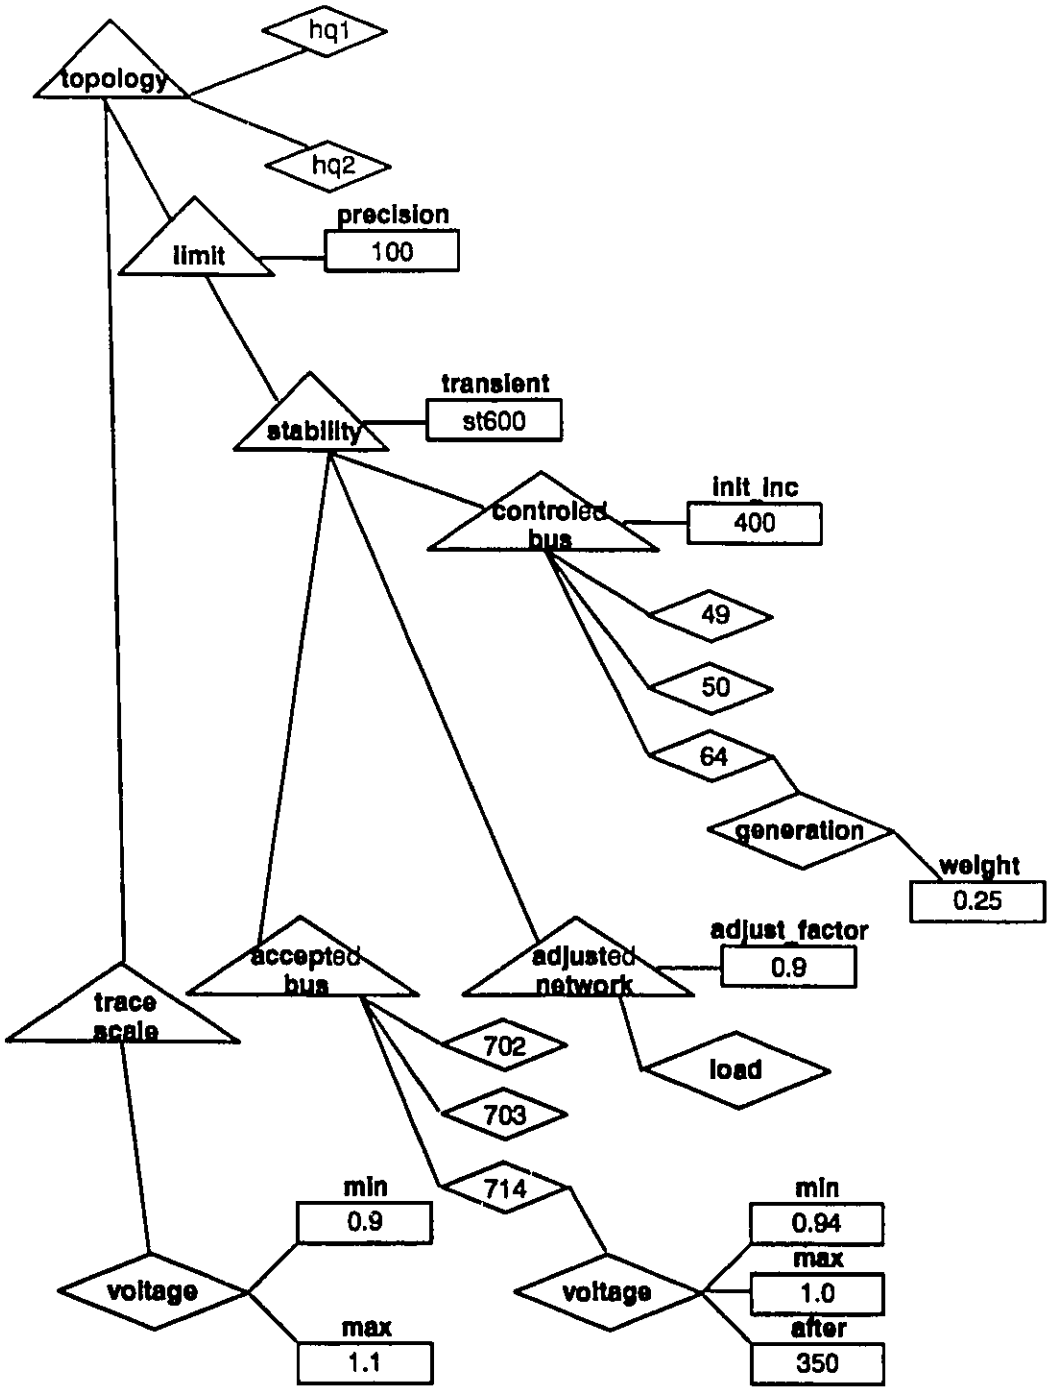

Fig. 3–10. Partial class–object network associated with script of Fig. 3–8.

its constituent classes and objects. The class–object network of Fig. 3–10 for the script of Fig. 3–8 shows that this is indeed the case. This class–object network is only partially complete for two reasons. First, **bus**es 49, 50, 702 and 703 are not fully defined: This was done for the sake of clarity. Second, many objects and properties are created in the course of the limit search process itself: These also have not been illustrated. For example, no details are given concerning **topology**s hq1 and hq2 and **bus**es 49, 50, 64, 702, 703 and 714 in relation to specific **stability**, **generation**, **load** and **voltage** objects generated as a result of performing different stability simulations. Nevertheless, it can be seen that class–object networks are useful in visualizing the evolution of a process in real time as it progresses towards its goal.

A particular feature of this class–object network is that of *upwards* inheritance: Every lower level quantity is required to characterize a higher level class or object. One also notices that the **control**ed, **adjust**ed and **accept**ed quantities are on an equal footing: Each is necessary to characterize the **limit** search under **transient stability**. In the same spirit, if a **simulate** process were included, the associated **control**, **adjust** and **accept** quantities would belong to a specific object of the **fault** class. The **trace** relates to the actions performed on each **topology** object and is therefore attached to this class.

Finally, the **init_inc** and **adjust_factor** properties respectively relate to the **controled bus** and **adjusted network** classes. This means that, in each case, the property is inherited by every object of these classes.

3.6 Conditional Process Control: Constructing Large-Scale Processes

The preceding sections have modeled and structured dynamic security analysis primarily in terms of relatively simple and straightforward processes. Typically, no more than a single rule structures the analysis and decision process in order to branch towards different outcomes. However, there is another kind of dynamic security analysis: The processes of the semantic net may themselves be viewed as elementary building blocks for constructing even more complex processes by means of conditional branching structures. In this section, we show how such large–scale processes can be built by means of two different semantic structures: The **while** and **if ... then ... else ...** structures. Because of the

(a) *The* **find** *is inside an* **if**

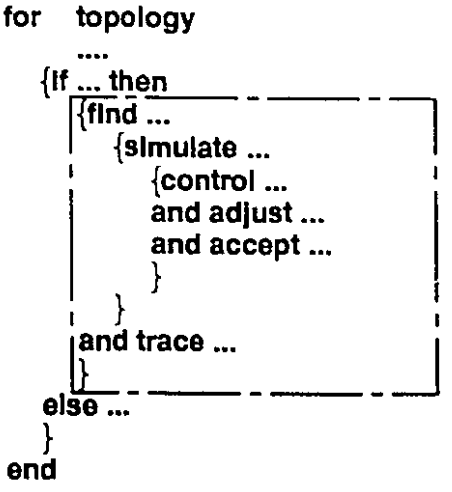

(b) *The* **find** *is inside a* **while**

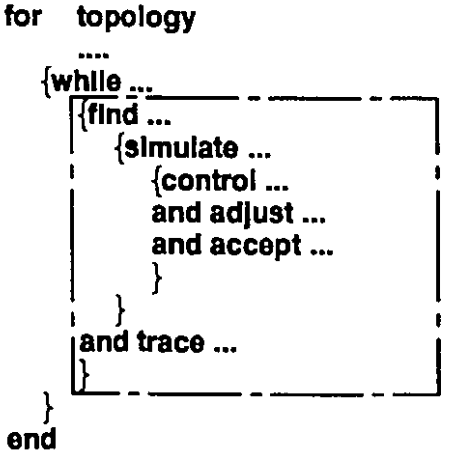

(c) *The* **find** *precedes an* **if**

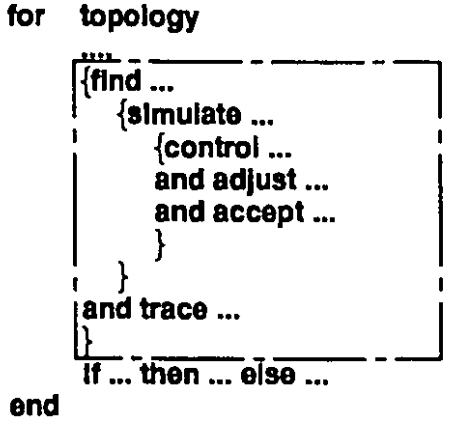

Fig. 3–11. The **find** process used with conditional branching structures.

semantic and syntactical complexity of such structures, for the sake of clarity, we will present examples in terms of scripts derived from the semantic net.

3.6.1 THE TWO TYPES OF CONDITIONAL BRANCHING

Large–scale dynamic security analysis processes can be built in two ways:

1) by using the **find** process as an elementary process within a conditional branching structure;
2) by building complex filters within an **accept** process.

Let us now consider each one in turn.

3.6.2 THE **find** AS AN ELEMENTARY PROCESS

Figure 3–11 illustrates the basic concept of considering the **find** as an elementary process. The **while** or **if** operate on a goal class which will eventually contain the result of the subsequent process. Where the targeted goal class has not be initialized, the default is to enter the **while** loop or the **then** portion of the **if**. Let us consider specific examples of how these language structures might be useful in practice.

Figure 3–12 illustrates a script where the objective is to find some combination of *additional* SVCs at buses 780 and 782 which will result in a security limit equal to or greater than 10 000 MW. This is done by means of an **if ... then ...** within a **for** loop. At the lowest level, the **find** first determines the security limit for the default data (i.e. the number of *additional* SVCs being 0), the contingency set being identified by means of the **simulate** process. If the security limit is at least equal to 10 000 MW, the process ends here. If it is less than 10 000 MW, the process increases the number of SVCs in the **bus set** by 1 SVC. This means that for every increment in the number of SVCs, *all* combinations of the *total* number of SVCs are generated by the 780,782 object and a security limit is determined for each particular combination (see Table 3–1). The process only stops when the desired security limit is attained or whether the **for SVC** loop continues to the end.

Figure 3–13 illustrates a script where the objective is to find a security limit using a different approach than that which first determines individual transfer limits for every contingency in a set, and then compares the limits (i.e. the

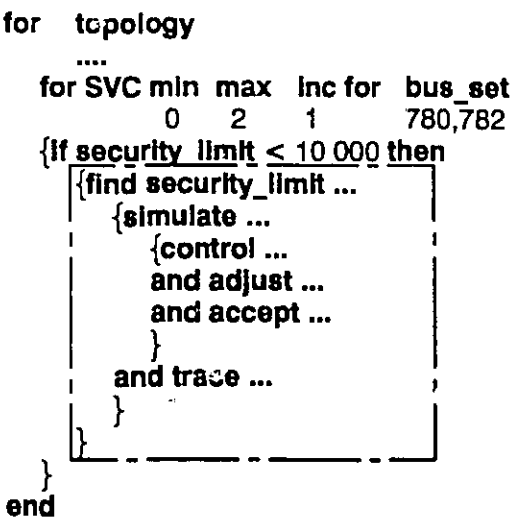

Fig. 3–12. Example of using the **find** process within a conditional branching structure: SVC configuration is modified incrementally to increase the security limit until the **while** loop is satisfied.

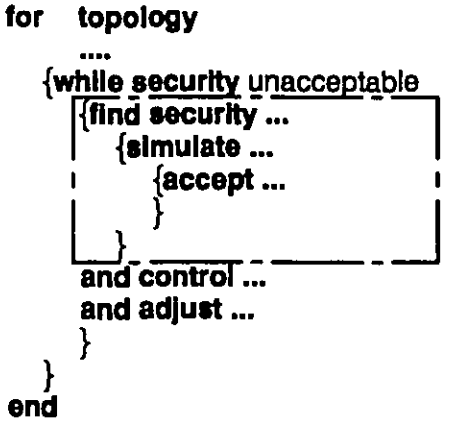

Fig. 3–13. Another example of using the **find** process within a conditional branching structure: An alternative strategy for security limit determination.

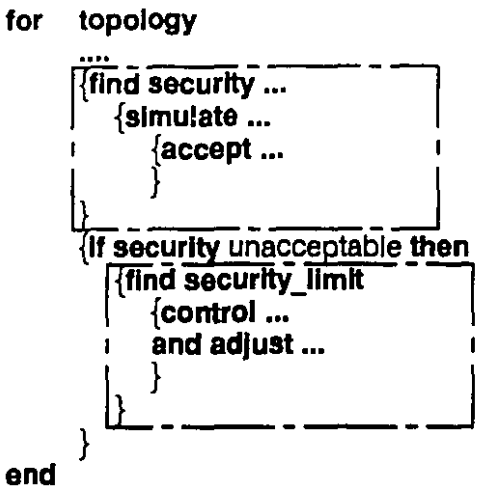

Fig. 3–14. Yet another example of using the **find** process within a conditional branching structure: A hybrid strategy for determining a security limit. The **If** follows and precedes a **find**.

default approach). This is accomplished by means of a **while** loop. We begin the search with a topology whose operating state results in acceptable cases for contingencies at certain locations and unacceptable cases at other locations. The script describes the following process: The **find** determines the **security** status for each of a set of contingencies, identified by means of the **simulate** process: Only the acceptance criterion need be specified. After having simulated the contingency at every location in the set, a test is performed on every object of the **security** class: Only those objects whose value is unacceptable are retained. The **control** and **adjust** processes are then applied to the latter, as a group and power transfer is reduced: This is because the **control** and **adjust** processes are not defined within the context of a **find**, and simply react appropriately to the **accept** analysis. Reentering the **while** loop, the **security** status of only the retained contingencies is found. This is repeated until no contingency location gives an unacceptable **security** status. When this is so, the last remaining case(s) with an acceptable **security** status determines the security limit.

Figure 3–14 illustrates yet another strategy for **security limit** determination. One might refer to this as a hybrid strategy. First, as before, the **find** determines the **security** status for each of a set of contingencies, identified by means

of the **simulate** process. Once again, only the acceptance criterion is specified. After having simulated a contingency at every location in the set, a test is performed on all objects of the **security** class by means of the **if**: Only those objects whose **security** is unacceptable are retained. Following this, the **then** specifies that a **security limit** is to be found among all remaining contingency locations by means of the default strategy: Consequently, a transfer limit is found for each of these remaining locations using a binary search, and the **security limit** is obtained by comparison of the limits obtained at these locations. Only **control** and **adjust** processes need be specified because the **accept** criterion is already known.

3.6.3 **accept** FILTERS

When determining transfer **limits**, complex criteria can be built within the framework of the **accept** process, leading to increases or decreases of the **control** variables within the context of a binary search strategy. In the absence of additional information, this search strategy is quite effective. However, if additional information can be obtained, either through complementary techniques or by extracting deep knowledge from a simulation, one can improve on binary search performance. One way of achieving this is through the use of more sophisticated filters.

One example of this is given in Fig. 3−15. This objective of this script is to find transfer limits under **long term stability**. This involves performing simulations on the order of 20 000 cycles long. Because these simulations are quite lengthy, it is useful to attempt to reduce their number. One method is to perform a post−contingency power flow before each long−term simulation. If the power flow is unacceptable, or does not converge, the full simulation need not be performed. If, however, the post−contingency power flow is indeed acceptable, then we have no choice but to execute the long−term simulation in order to assess power system voltage stability.

Another example is shown in Fig. 3−16. This example anticipates a result that is obtained in Chapter 5 of the present thesis. For the moment, we describe the filter and refer the reader to 5.6.5.1. In this example, when at least two stable simulations have been performed, the **pu_signal_energy_slope** (i.e. per unit signal energy slope) is calculated. If the slope is greater than some value (in this

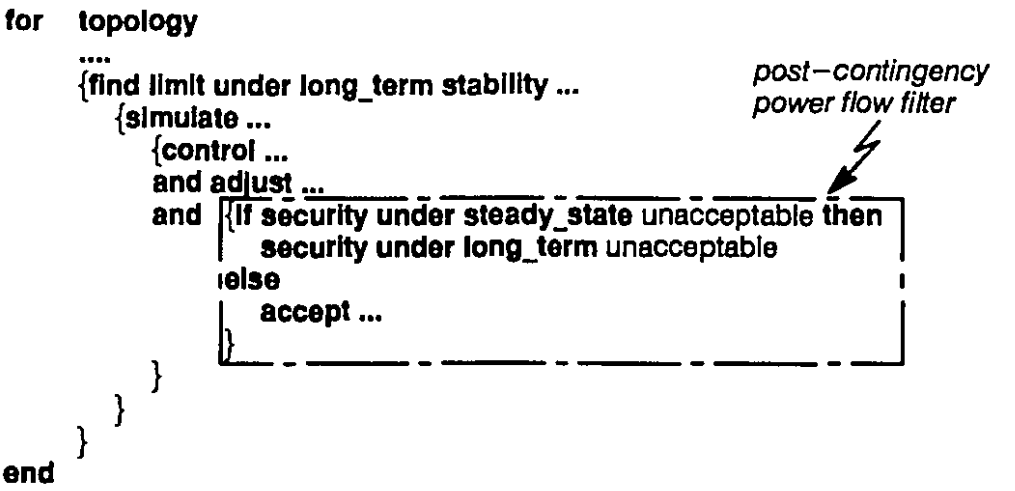

Fig. 3−15. Example of the filter concept: A post−contingency power flow precedes long−term stability simulations.

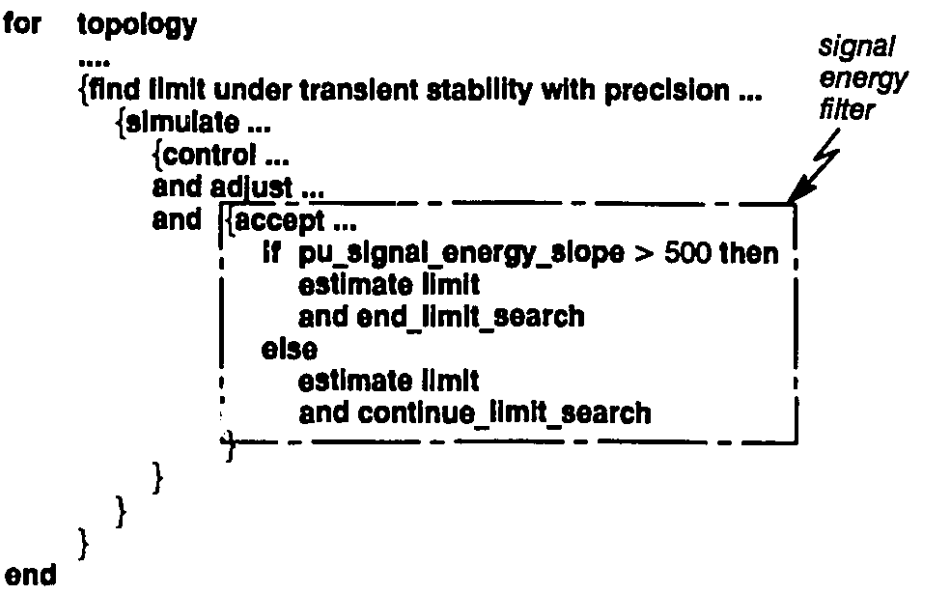

Fig. 3−16. Another example of the filter concept: Using signal energy behaviour to accelerate stability limit determination.

case, 500), a stability **limit** is **estimate**d and the limit search need not continue as the accuracy of this estimate is known to be high. If the **pu_signal energy_slope** is not high enough, this estimated case is the next case to be simulated: If the case is stable, a new stability limit is estimated, the per unit signal energy slope is calculated anew, and the search comes to a halt if the slope is in the correct range. If it is unstable and the distance to the closest acceptable case is less than the required **precison**, the search also comes to a halt. If it is unstable and the distance to the closest acceptable case is greater than the required **precison**, a point is chosen halfway between the two and the process is repeated until a third stable simulation and an appropriate p.u. signal energy slope are obtained. If the base case represents a typical system operating point (i.e. peak load), the limit can be found with high accuracy from at most 3 stable simulations.

3.7 SUMMARY

In the present chapter, we presented a taxonomy of the *processes* of dynamic security analysis. To build such a taxonomy, it was found necessary to combine semantic nets, classes, objects, properties and rules, thereby permitting the full spectrum of knowledge in dynamic security analysis to be represented. The semantic net used to build the taxonomy is essentially structured around the goals of dynamic security analysis, and the processes are found to be highly goal–dependent.

Semantic nets are not only a useful means of structuring dynamic security analysis, they can also be *interpreted* as describing the language of operations planners. Descriptions of dynamic security analysis processes can consequently be written down precisely, and many examples were presented of scripts where powerful concepts were represented with clarity and concision. The semantic net is also seen as defining hierarchical relationships between classes and objects in the net, and class–object networks can be constructed on this basis. This is particularly useful in visualizing the different properties and values of a specific process, particularly as it progresses towards its goal.

Though the taxonomy itself is limited to relatively simple and straightforward processes, its component processes can be viewed as building blocks for describing complex, large–scale processes when combined with conditional

branching, filters and looping.

Because the resulting language has semantic *and* software generality in that the processes are defined independently of any commercial simulation software, this language can be used as a basis for constructing a very high–level *shell* or *framework.* Unique bridges need only be built for each application program supported in a software *library*. A particular *framework* implementation is considered in the next chapter.

CHAPTER 4

THE ELISA FRAMEWORK PROTOTYPE: DESCRIPTION AND PERFORMANCE

4.1 INTRODUCTION

As we have seen in the previous chapter, the object of software *frameworks* is to mechanize high–level tasks which call upon diverse software tools and which are normally performed by experts. The present chapter focuses on the ELISA framework prototype which performs power–system dynamic security analysis in the operations planning environment. ELISA is an acronym for *Estimateur de LImites de Stabilité Automatisé* which means "Automated Stability Limit Estimator". ELISA mechanizes routine but complex processes traditionally carried out by experts. These processes are essential to power–system dynamic security analysis and their mechanization greatly accelerates their realization. Typically, ELISA executes appropriate power–flow and transient–stability simulations (i.e. using commercially available simulation software), performs result analysis, identifies and executes changes to the input and repeats this process until a user–defined goal, such as finding transient stability transfer limits, has been achieved.

The user describes dynamic security analysis processes by means of the object–oriented language formally presented in the previous chapter. The semantic net used to define this language not only defines the particular aspects implemented within the ELISA prototype (i.e. the shaded portions) but constitutes a roadmap to future ELISA enhancements.

Three different versions of the ELISA prototype have extensively been tested and used by the Hydro–Québec operations planning staff since November 1991. To illustrate the benefits of frameworks and the flexibility afforded by the ELISA prototype, a number of examples of ELISA scripts are presented. These include large–scale sensitivity studies performed on the Hydro–Québec network in a fraction of the time formely required. It is seen that the prototype reduces study–cycle time, permits the execution of very large and ambitious studies which normally would not have been performed for want of resources

and, indeed, enables operations and system planners to work at a more strategic level.

4.2 FUNCTIONAL OVERVIEW

4.2.1 ELISA FEATURES

The ELISA framework prototype is a script–driven shell which reads a natural–language–like script which in turn defines the process goal and describes the tasks to be performed. Figs. 3–8 and 3–9 are typical of scripts read and executed by ELISA.

Figure 4–1 shows how the syntax of Figs. 3–1 to 3–5 is used to build ELISA scripts. An arbitrary number of topologies is entered under the `topology` class for any given session. ELISA presently determines a single dynamic security analysis goal, that of `transient stability` transfer `limit` determination. This includes long–term stability evaluation provided that the transient stability software supports transformer tap changer and shunt reactor disconnect control system models. ELISA determines this goal within a user–defined `precision`, using the modified binary search described in 3.3.2. Only one particular `transient stability` software is presently supported to perform the search: This is the ST600 commercial transient/long–term stability package [Valette, Lafrance, Lefebvre & Radakovitz 1987]. Power flows are performed using the RP600 power flow program [Gaba, Audette, Guillemette & Lafrance 1993] which is compatible with the ST600 software. Though these are the only software presently supported, ELISA is not inherently tool–dependent: Bridges to other software packages can also be realized.

The user has the option of requesting the calculation of the correct SLG fault impedance for the contingency present in the topology data file (i.e. `determine SLG fault impedance`). Alternatively, the user can request a contingency defined directly within the script (i.e. using `simulate for fault...`). Of course, the user can choose to use neither option.

The user can define process rules whereby `generation` or `load` is `controled` at `buses`, `zones` or on the whole `network`, and `generation` or `load` can be `adjusted` at `buses`, `zones` or on the whole `network`. In the present ELISA prototype, either `buses` or `zones` can be individually specified, but not

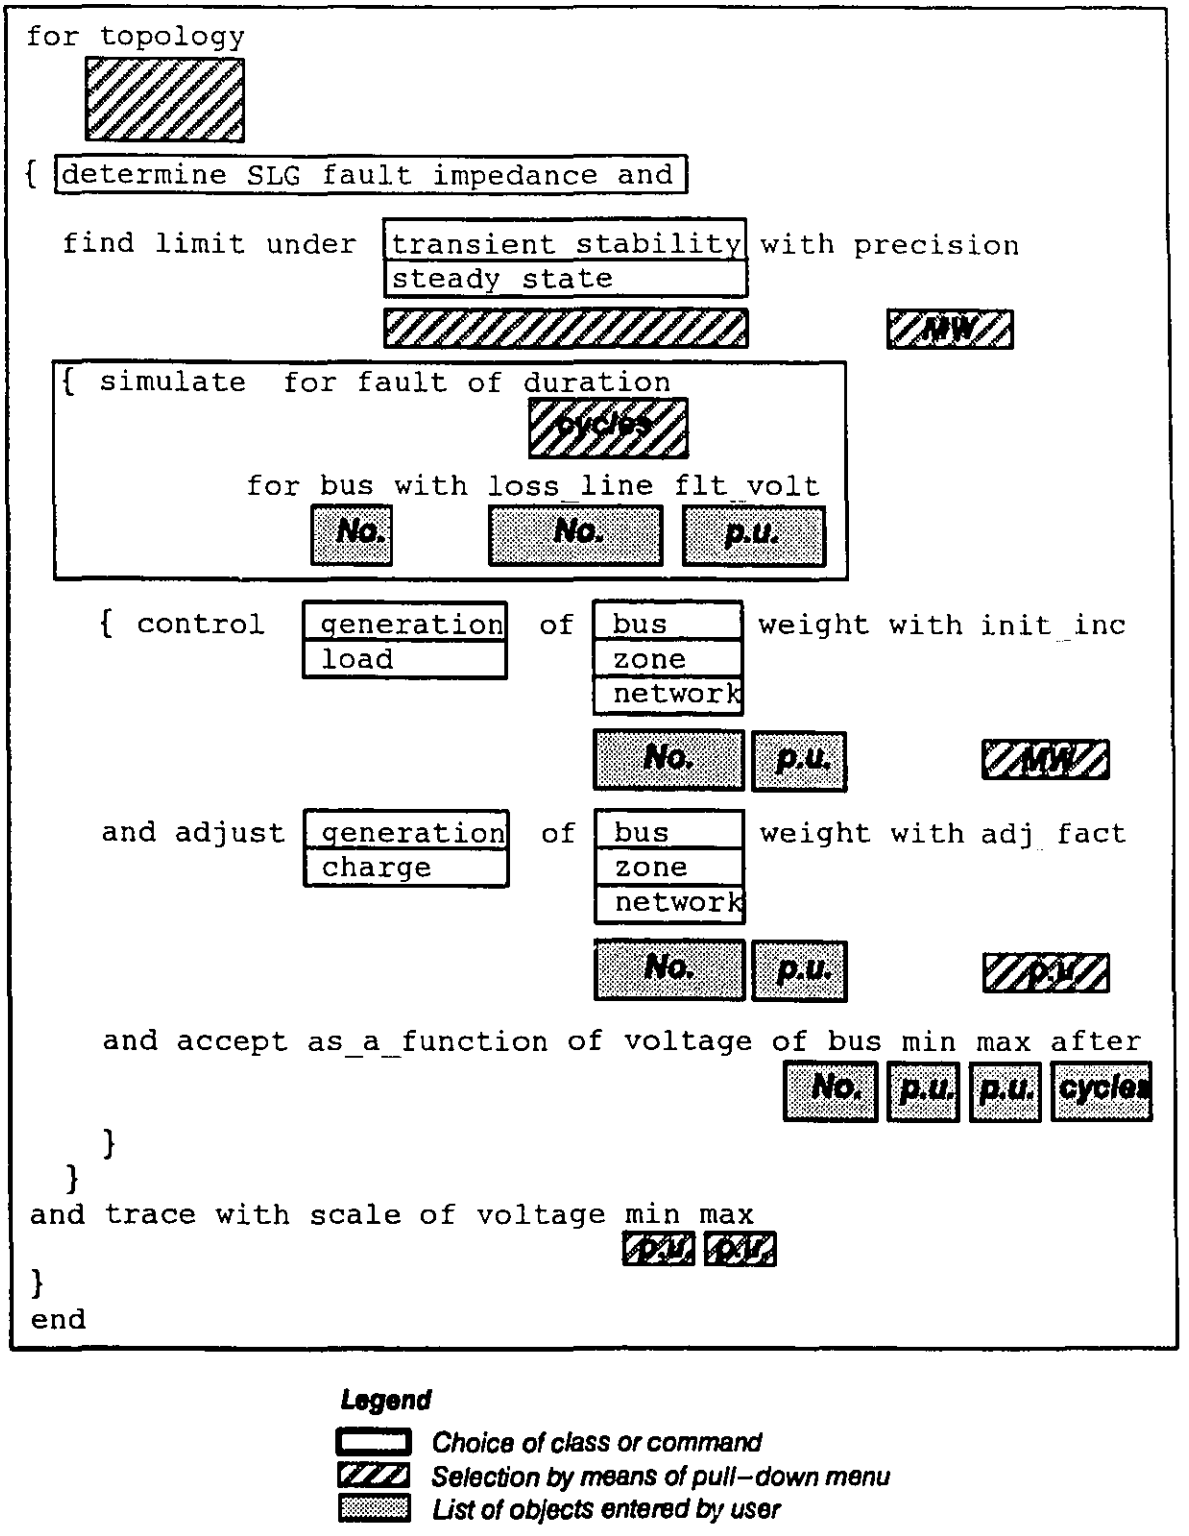

Fig. 4–1. Creation of ELISA scripts for transient and long–term stability limit searches.

both. However, there is no limit to the number that can be specified.

Individual cases are screened on the basis of `voltage` where the `minimum` and `maximum` are identified at each monitored `bus` after a specific `time`. In addition to this, there is an implicit criterion which infers that any simulation which ends prematurely is considered unacceptable: A detailed explanation of this is given in 6.3.2.2.

The user also has the option of plotting output results on a pre−selected scale (i.e. `trace ...`). At the end of every limit search, ELISA plots the voltage of the monitored bus for every case generated by the limit search, as in Fig. 2−2. If the criterion includes many buses, the results of each bus are plotted in turn. Also at the end of the search, ELISA automatically draws the network single line diagram showing the power flows for the case corresponding to the limit.

The various combinations of all of these features permit the preparation of over one hundred different script types within the present ELISA prototype.

4.2.2 ELISA INTERFACES

There are two ways of creating ELISA scripts. The user simply enters a line or screen editor, creates a file in the format shown (Fig. 4−1), returns to the UNIX environment and executes this script by issuing the command:

`csh.elisa` ***script_name***

An alternative means of building scripts is by means of a custom−built, Windows interface for ELISA, developed by Hydro−Québec under the SUN Open Windows environment. This is the preferred approach for most operations planners: Pull−down menus are liberally provided which permit a user to create scripts essentially by pointing and clicking on the screen. A minimum of keyboard input is required, as can be seen from the legend in Fig. 4−1 for the transient/long−term stability interface. A primitive form of data validation is performed which ensures that only integer or real number input of the appropriate length is entered in the appropriate fields. Otherwise, pull−down menus constrain users to only valid choices.

Principal directory:
`liv_92.3`

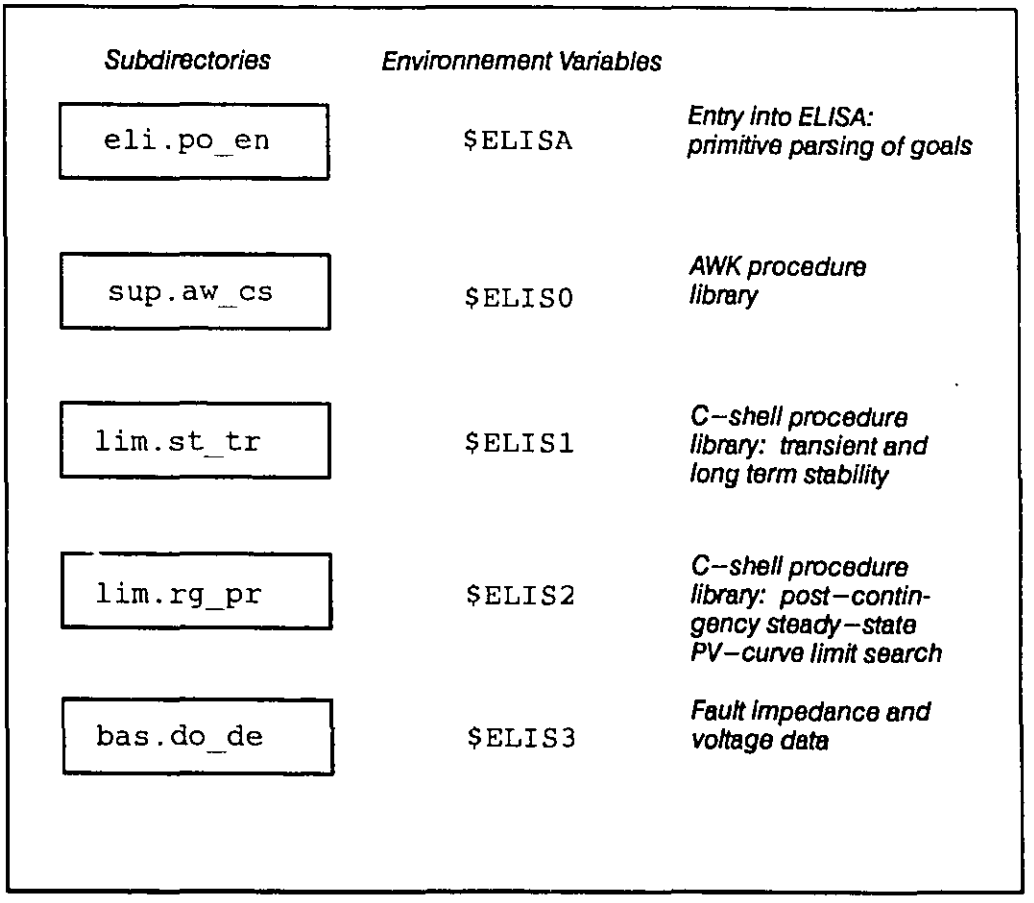

Fig. 4–2. ELISA directories and environment variables.

4.2.3 ELISA AND OPERATIONS PLANNING METHODOLOGY

As we shall see shortly in the examples, in mechanizing expert processes, frameworks have the potential of increasing the productivity of individuals to an unheard–of degree. The brunt of planning activities will shift from performing routine tasks to accomplishing far more strategic functions.

One such function is data verification. As pointed out in 2.4.6, it is important to calibrate simulations with known network response: Planners will have more time to attend to such crucial activities. Another aspect of this is that base case power flow and stability simulations can be prepared more carefully, with greater attention being given to such oft–neglected aspects as load distribution, load characteristics and transient load behaviour. Additionally, more time is available for identifying a greater number of umbrella topologies: This will ensure that security limits will be less conservative for lightly degraded topologies (see 2.4.4). Finally, more sensitivity studies will be performed on voltage support equipment distribution (i.e. SCs, SVCs, etc.), thereby facilitating the preparation of detailed security control strategies (i.e. limit tables, etc.).

A few points should also be stressed regarding operations planning methodology. As recommended in 2.4.7, before ELISA is used, it is essential that the base case power flow include fictitious generators on every EHV bus: This not only permits voltage profile to be excluded as a factor in subsequent transfer limit analysis, it ensures one–shot convergence of the power flow at every step of the limit search. When these are transferred to the stability software, the corresponding reactive power flows are converted to passive shunt impedances. This was referred to as a continuous rather than a discrete approach to positioning reactive power reserves on the network, subject to *a posteriori* verification. In the absence of more effective standard power flow software or the use of OPF, this constitutes a *prerequisite* to the effective use of ELISA.

Finally, in chapter 2, we saw that the dispatching of multiple generating units in power stations can effectively be treated as a single generator with continuously variable machine and transformer impedances (see 2.4.3). This method is recommended in view of its simplicity and of the conservative limit values that are derived therefrom. The ST600 software was modified by Hydro–Québec, in parallel with the development of ELISA, to permit the

choice of this method of treating multiple generator power stations; the RP600 software was simultaneously modified to do the same for generator transformer impedances: Consequently, these features should always be used. For simulation software not equipped with this capability, it will eventually be necessary to include these methodologies directly within ELISA.

4.3 SOFTWARE ISSUES

4.3.1 SOFTWARE DEVELOPMENT

ELISA is a library of programs written in the UNIX C–shell [Anderson & Anderson 1986] and AWK [Aho, Kerhighan & Weinberger 1988] programming languages. A breakdown of the contents of its various directories is provided in Fig. 4–2. There is a very clear division of tasks between the two programming languages, as we shall see.

The UNIX C–shell is a language which enables the programming of operating system commands, and their subsequent execution, by means of an interpreter. The significance of this is considerable: Any software ported to the UNIX environment can therefore be incorporated within a UNIX program. From a strategic perspective, the capacity to *program* an operating system therefore opens the door to the creation of *higher–level software*, including database technologies, by means of operating system language scripts.

In ELISA, the basic task of C–shell scripts is to identify the goal, enter the appropriate process, choose the right application software, prepare the data and scripts required to run this software (i.e. in the software's own extension language), perform high–level tests based on the analysis of results, determine whether or not the goal has been achieved, and reenter the process at the appropriate point if this is not the case.

Though originally conceived as a UNIX utility, the AWK language has evolved to the point that it is now a complete, interpreter–based, C–like programming language, capable of performing rapid pattern–matching in multiple files from a minimum of coded instructions. It is also capable of writing or adding information to files, and has many of the features of C, which makes it particularly useful for result analysis.

The necessity of using AWK stems from the fact that C–shell scripts are incapable of reading the contents of files and thus perform direct result analysis. Let us explain the importance of this. Large simulation software such as the RP600 and ST600 power flow and transient/long–term stability programs often save their results in binary files which can only be accessed usefully by means of resident software features or post–processors. The extraction of data for result analysis therefore requires returning to this software by creating unique scripts for accessing these data files, and causing data to be written in ASCII format in yet other files. C–shell scripts can be used to mechanize such procedures. However, AWK scripts are required to read the contents of the ASCII files, perform result analysis, and either write this result in temporary files or store them in C–shell variables. This explains the use of the C–shell/AWK tandem to create the basic ELISA framework infrastructure.

4.3.2 PHILOSOPHICAL CONSIDERATIONS

In the light of the previous discussion, a certain number of ideas concerning operating systems merit further exploration. Any software that can be executed under a particular operating system can be seen as *de–facto* integrated within this operating system. Programmable operating systems can be seen as going a step beyond, providing the means for creating ever higher–level software from elementary building blocks which are neither elementary nor meant to be used as "building blocks". Though expert–system shell technologies and artificial intelligence (AI) languages such as LISP and PROLOG attempt to provide environments for creating high–level strategies and interfacing with a variety of software including databases, such tools are inherently limited in that they are themselves resident within and dependent on operating systems, and are not *in total control* of their environment. The operating system sits alone at the top of the software control pyramid.

In the long term, the quest for machine intelligence must strive towards total control of the software environment for one very simple reason: Intelligence is the highest form of software and intelligent software, by definition, can only reside at the top of the software pyramid. Consequently, operating systems technology must inevitably converge with knowledge representation and processing technologies. In the short term, a natural evolution of operating systems would be to offer object–oriented modelling and programming features,

including the capability to build rule bases, knowledge bases, databases, and to employ inference engines.

The strategy of controling one's environment invites additional comment. We have seen in Chapter 1 that a fundamental tenet of security (and survival) is to exercise the greatest possible control over one's environment. Let us now reason backwards: If the known strategy is to exercise control over one's environment, one infers that survival is a key goal. Transposing this to the sofware environment, if control over the environment constitutes a necessary condition for the emergence of intelligence, survival is therefore an inherent trait, or goal, of this intelligence. One therefore reasons that the quest for survival is a *prerequisite* to intelligence, or alternatively, that intelligence arises as a *by–product* of some demonstrated capability to survive.

In a recent interview given by Allen Newell, an acknowledged pioneer in the field of artificial intelligence, one sees that these conclusions parallel many of his own ideas [Chandrasekaran 1993]. For example, Newell states that "learning [is] a side effect, learning [is] something that [goes] on continuously and not at the volition of the agent". If the capacity to learn is considered such an essential characteristic or feature of intelligence that the two are virtually indistinguishable, this statement can be considered identical to the conclusion of the preceding paragraph. Newell also states that "intelligence is not just computation, it is related to the achieving of a goal ... It relates to whether you can get at the proposition in your knowledge base when you need it for your goal". In the preceding paragraphs, we have gone a step further and proposed that the fundamental goal of intelligence is survival. One might say that survival is the goal which enables intelligence to establish priorities in selecting other goals.

All of this might seem to diverge from the main topic were it not for the fact that a primary concern of this thesis is the dynamic *security* of power systems. We have already defined power system security as the art and science of power system *survival*. Pursuing our previous line of reasoning, one can view the power system security problem as an ideal problem in artificial intelligence. A software capable of assimilating its own survival to that of the power system, and incorporating all of the necessary attributes to ensure its survival (i.e. real–time monitoring, high–level reasoning, learning capability, access to simulation

tools, ability to analyse results, ability to execise control on its own computer environment and on the network, etc.) may well constitute an important research direction towards achieving a software–based intelligence ultimately greater than the sum of its parts.

4.4 TRANSIENT AND LONG–TERM STABILITY: INTERNAL PROCESS FLOW

Figures 4–3 to 4–7 represent a simplified flow chart of the main ELISA internal process flow for processes relating to transient and long–term stability. The present section explains the main features presented in these figures. An important remark: ELISA neither compiles nor interprets its scripts in a rigourous sense. ELISA scans the script and extracts the information that it needs as it accomplishes intermediate tasks and progresses towards the specified goal.

4.4.1 ELISA: THE KEY DIAGRAM

Figure 4–3 presents the key ELISA flow chart. The names of C–shell procedures always begin with the suffix `csh.`; accordingly, AWK procedures are identified by the suffix `nawk.` which stands for "new AWK", a more recent version of the AWK language. Let us now examine some of the details involved with these procedures.

The `csh.elisa` procedure is the entrance to ELISA. Here, the goal is identified and the appropriate procedure is called. If the goal is to find a transient or long–term stability limit, `csh.limite` is called. If the goal is to find a steady–state post–contingency PV–curve limit, the `csh.pc_lim` procedure is called (see section 4.5).

The `csh.limite` procedure incorporates the highest–level loop: The topology loop. The topologies are counted and identified, and the remainder of the ELISA script is executed for each one in turn. Before a stability limit search is initiated, this procedure verifies the existence of the appropriate power flow and stability data files (i.e. the `so.`*topology* and `ma.`*topology* files, respectively) which are of course required. ELISA recognizes power flow data files by the suffix `so.` and stability data files by the suffix `ma.`; *topology* is a unique name which identifies the topology for which a limit is to be found (including the associated stability data and default contingency). Having looped through the last topology, a clean–up procedure eliminates all unnecessary temporary files. The `csh.pc_lim` procedure performs a similar function for the steady–state

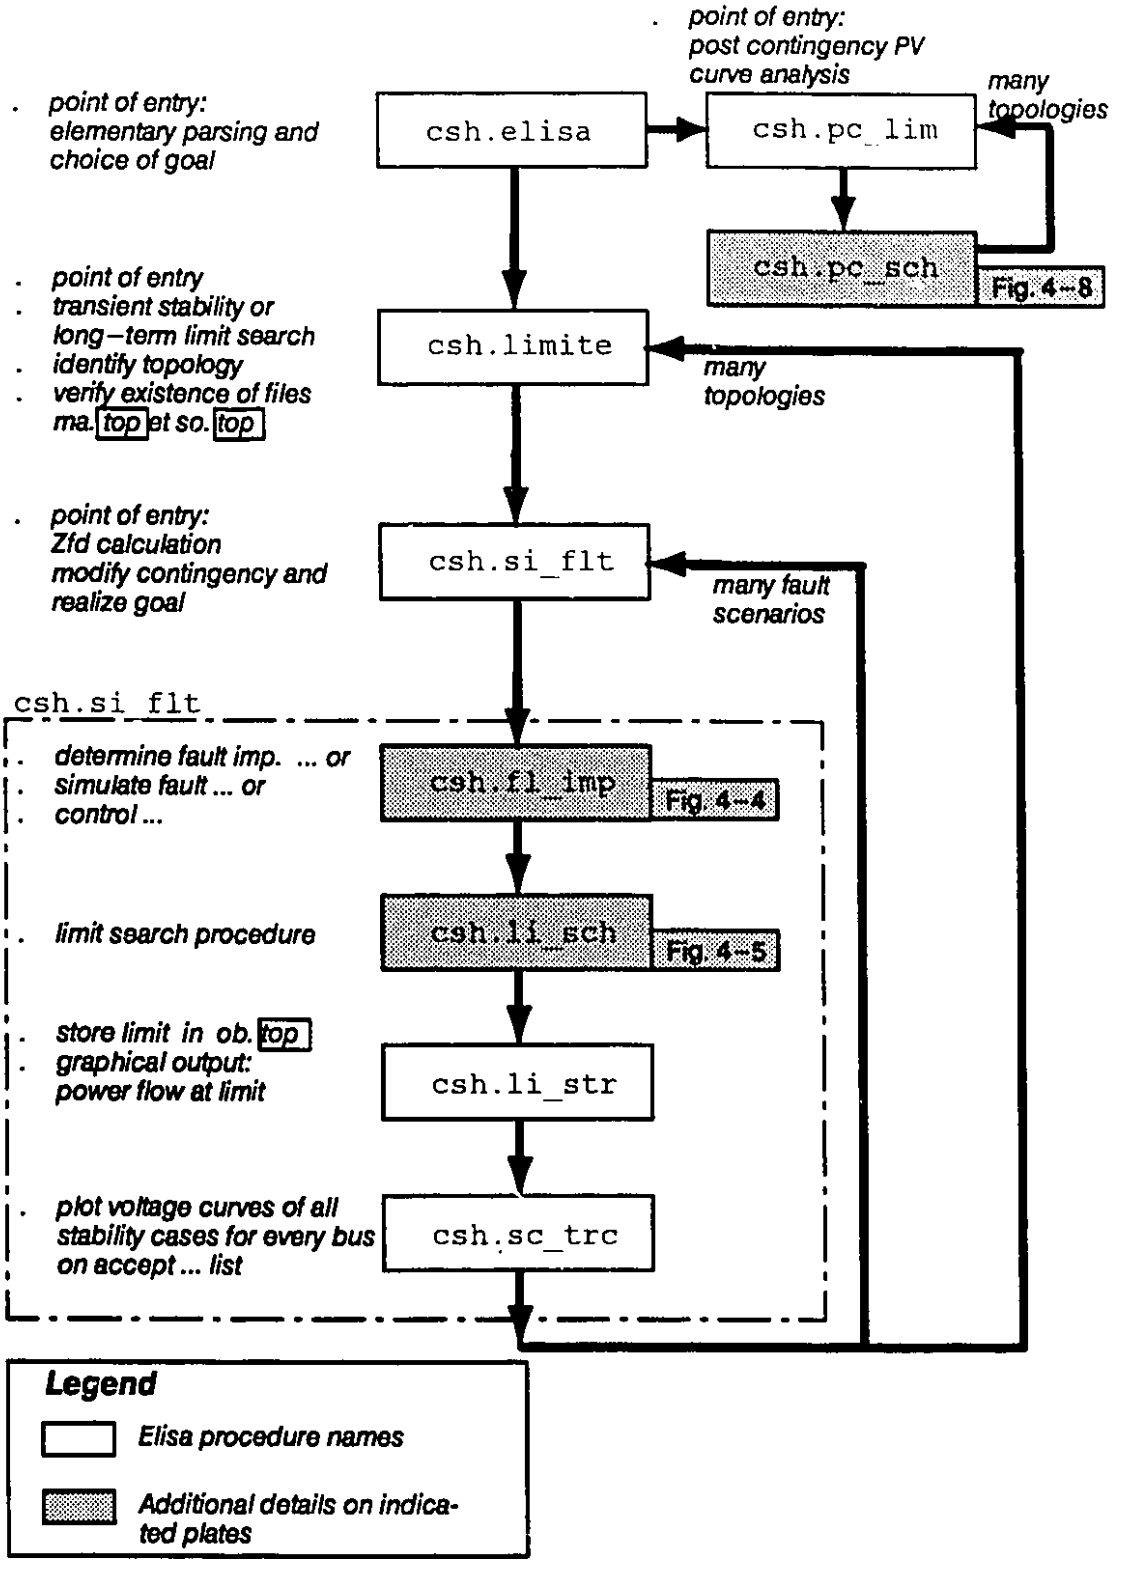

Fig. 4-3. ELISA procedures and logical flow: Key diagram.

post−contingency PV−curve limit search process except that the existence of the appropriate power flow and contingency files is verified at the outset (i.e. respectively the `so.`*topology* and `co.`*topology* files).

The `csh.si_flt` procedure is then called. This procedure incorporates the looping code required to execute limit searches for successive fault scenarios, as shown in the previous chapter (Figs 3−2 and 3−9). If this is not required, SLG fault impedance determination is initiated (i.e. `csh.fl_imp`). Upon completion of this task, or if it has not been requested, the limit search procedure is initiated (`csh.li_sch`). When this procedure has completed its task, the limit is extracted from individual simulation summary data files (i.e. the `ob.`*topology_case* files relating to individual simulations or cases) and stored in the limit−search summary file for the topology (i.e. the `ob.`*topology* file). The voltage curves are then plotted as in Fig. 2−2 for every bus identified within the acceptability criterion of the ELISA script, and a single line diagram of the power flow is drawn automatically for the limiting case (i.e by `csh.sc_trc`).

4.4.2 DETERMINING SLG FAULT IMPEDANCE: THE `csh.fl_imp` PROCEDURE

As shown in Fig. 4−4, if the `csh.fl_imp` procedure is called in the context of a `simulate` command, a *primitive* stability data file is first generated: This primitive stability data file consists of the original file *excluding* whatever contingency data which may be present, and this is stored under the filename `ma.`*topology*`_pr`. A new stability data file is then created using the contingency requirements defined in the ELISA script, and this file is created under the `ma.`*topology_dur_bus* filename format where *topology* is the original topology name (i.e. found in the ELISA script), *dur* is the value of the fault duration found in the script, and *bus* is the fault bus identification, also found in the script. For ELISA, the *new* topology name is now *topology_dur_bus*.

If the `csh.fl_imp` procedure is called in the context of a `determine fault impedance` command, the procedure enters at this point. Otherwise, it continues by first copying the valid stability data file, now containing contingency data, having the form `ma.`*topology*, to a filename in the form `ma.`*topology*`_or`. The simulation time is now modified to 0.6 cycle, the time−step is set to 0.2 cycle and the test impedance is read from an internal ELISA data file named `ob.fault_dat`.

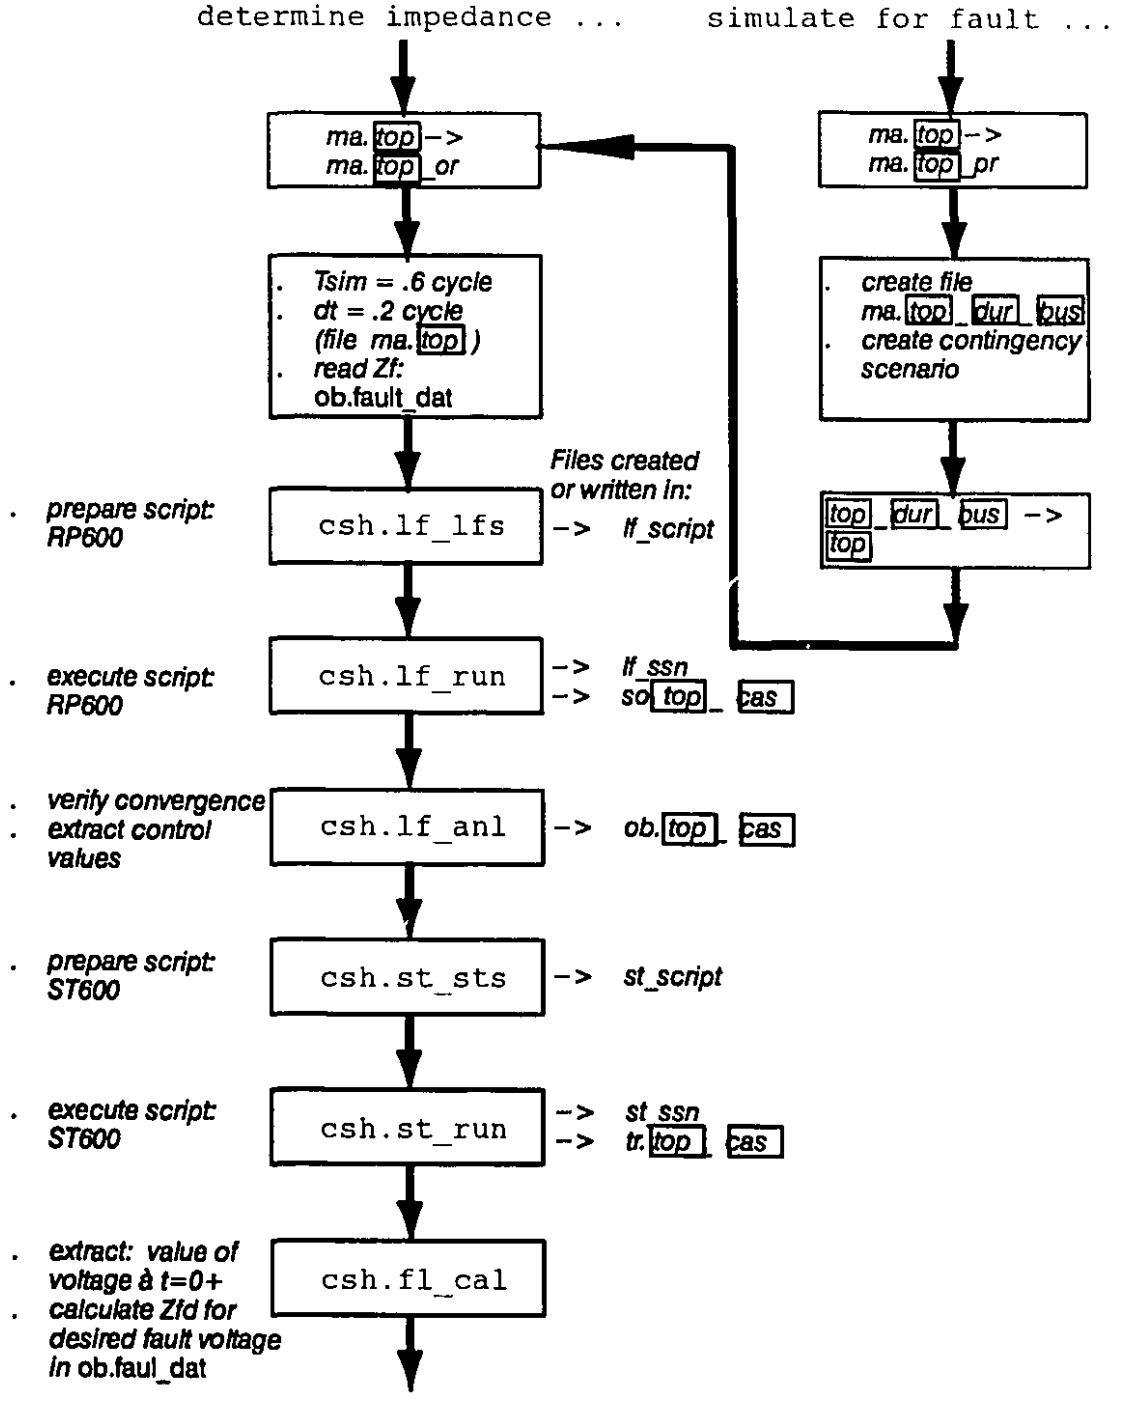

Fig. 4–4. ELISA process flow: The `csh.fl_imp` procedure for fault impedance computation.

Then, the csh.lf_lfs procedure prepares the script required to execute the RP600 power flow software, and stores it in a file named lf_script. This script also includes the name of the binary file in which results are to be stored (i.e. in the form so.*topology_case*). One might say that the csh.lf_lfs procedure generates, on its own, a program in the RP600 language. As the process goal is to calculate the fault impedance, the *case* name is impedance.

The csh.lf_run procedure is then called, executing the RP600 power flow according to the instructions contained in the script. The interactive power flow session, created in the course of the program's execution, is redirected and temporarily stored in the lf_ssn (ASCII) file. This file also contains the values of the control and adjusted variables as these have also requested in the RP600 script. The csh.lf_anl procedure analyses these results, first verifying for convergence of the power flow (this information is stated explicitly in the lf_ssn session file), then extracts control and adjust variable values and stores these in the ob.*topology_case* file.

The csh.st_sts procedure then generates the script required by the ST600 stability software, and stores it in a file named st_script. This file also contains the name of the binary file in which results are be stored (i.e. in the form tr.*topology_case*). The csh.st_run software then executes the script, the results are stored in binary form, and the interactive stability session is redirected and temporarily stored in the st_ssn (ASCII) file. The csh.fl_cal procedure then uses the STPP post-processor (associated with the ST600 stability package) to extract the value of voltage at t = 0.2 cycles (i.e. at t=0+) from the stability binary result file, permitting the calculation of the correct fault impedance as outlined in 2.4.8. This is accomplished essentially as outlined in 4.4.5 in the case of the csh.st_anl procedure. The value of the impedance is stored in the topology limit search summary file, ob.*topology*.

When the process is terminated, all files under the impedance case name are discarded, and all of the lf..., st... and ta... ASCII files are no longer required and can be written over.

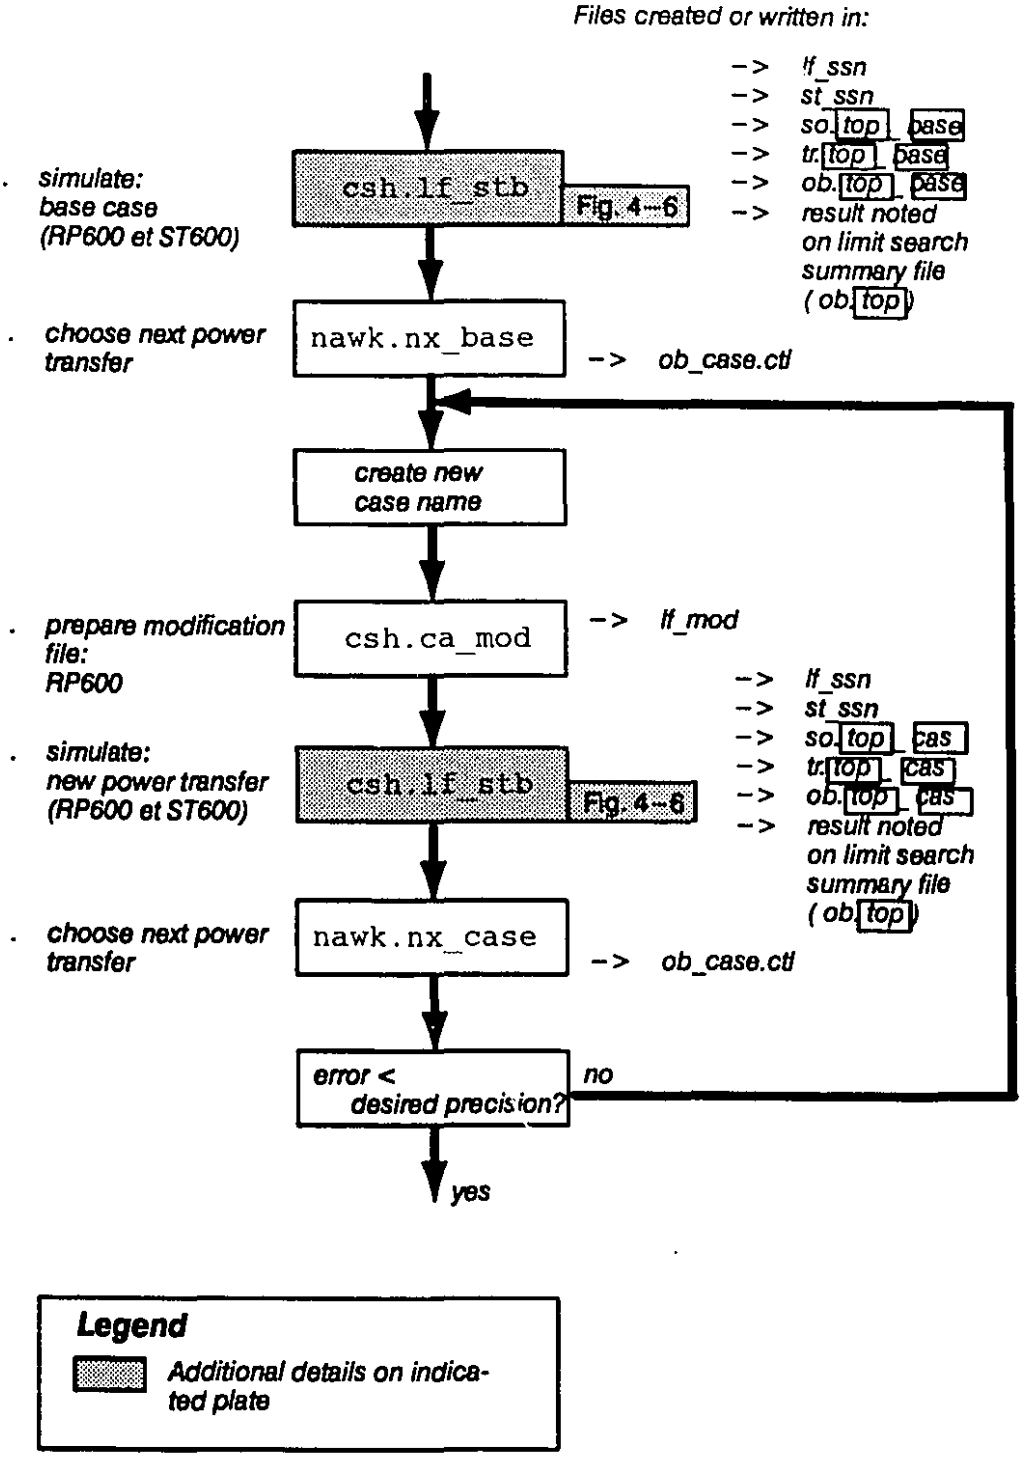

Fig. 4–5. ELISA process flow: The `csh.li sch` stability limit search procedure.

4.4.3 DETERMINING TRANSFER LIMITS: THE csh.li_sch PROCEDURE

The csh.li_sch procedure, illustrated in Fig. 4–5, is an extremely important ELISA function. First, the csh.lf_stb procedure is called: This executes the base case power flow and stability simulations exactly as described above, except that the *case* name is now base rather than impedance. The acceptability of the simulation is noted in the topology limit search summary file, ob.*topology* (and on the user's screen). The csh.lf_stb procedure is treated in greater detail in Fig. 4–6.

The nawk.nx_base procedure reads the base case result, stored in the ob.*topology* file, and determines the next power transfer using the value for the initial increment, init_inc, found in the ELISA script. This information is used to create a new *case* name which is used to identify subsequent power flow and stability results. For example, if the power increases by 400 MW, the new *case* name is p400; subsequently, the binary power flow result file will have the name so.*topology*_p400 and the binary stability result file will have the name tr.*topology*_p400. Of course, if power is decreased by 400 MW, the case name becomes m400.

The csh.ca_mod procedure modifies the input data for the next power flow case and the csh.lf_stb procedure is again executed, this time with the new *case* name. The nawk.nx_case procedure then determines what the next case will be using the modified binary search strategy (see 3.3.2), and the process continues as shown until the limit error corresponds to the desired precision.

4.4.4 EXECUTING POWER FLOW AND STABILITY SIMULATIONS: THE csh.lf_stb PROCEDURE

This procedure is illustrated in Fig. 4–6. We are already familiar with this procedure as its first five components have been described at length in 4.4.2. The only difference here is that an appropriate case name has been created in the calling procedure to reflect the particular case to be simulated. Afterwards, the csh.st_anl procedure extracts the desired voltage waveform data from the stability binary data result file and performs acceptability analysis based on the guidelines provided in the ELISA script. As before, at the end of this procedure, the lf..., st... and ta... ASCII files are no longer required and can be written over.

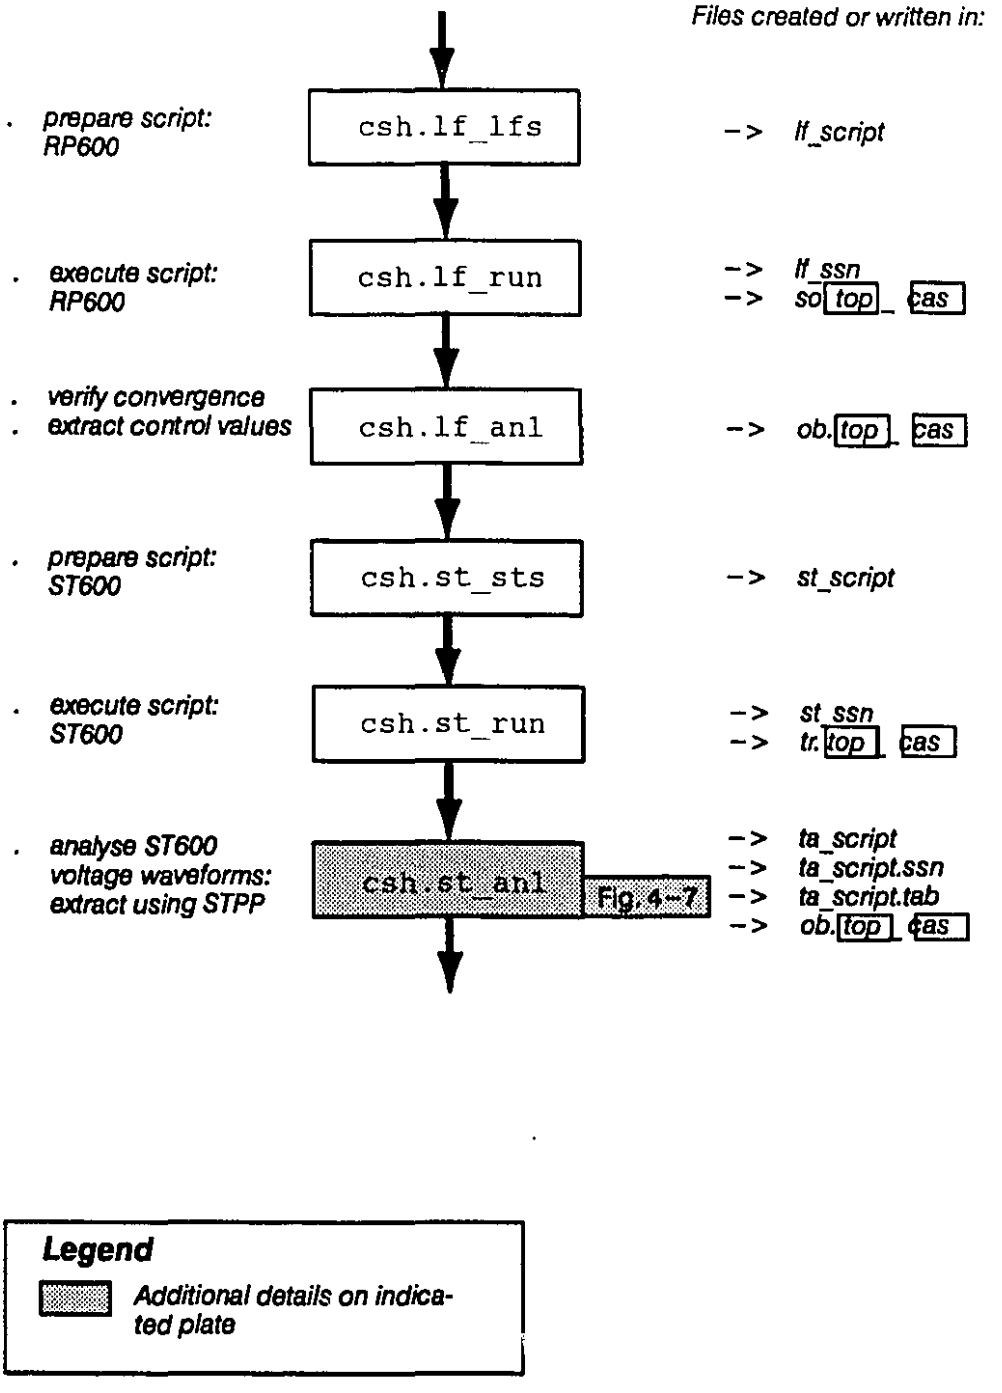

Fig. 4–6. ELISA process flow: The `csh.lf_stb` power flow and stability execution procedure.

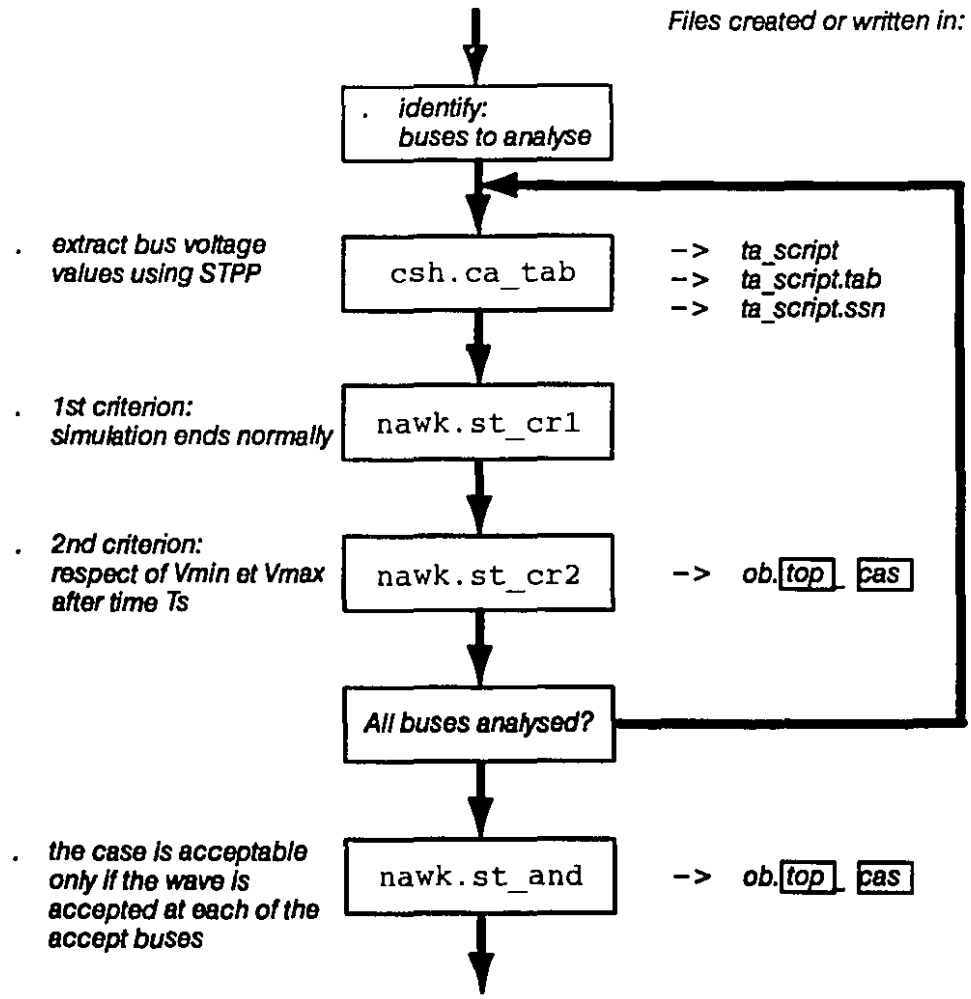

Fig. 4–7. ELISA process flow: The `csh.st_anl` stability result analysis procedure.

4.4.5 Analysing stability simulation output: The `csh.st_anl` procedure

This procedure is summarized in Fig. 4–7. First the number and identities of the monitored buses are read. Then, each bus is treated as follows, one at a time:

The `csh.ca_tab` procedure extracts bus voltage values as a function of time from the stability binary result file using the STPP post processor and copies these into a temporary file named `ta_script.tab` in ASCII format. To do this, it first writes a script in the language of the STPP binary data

post–processor (i.e. `ta_script`), and executes this script using STPP. The interactive session is redirected to yet another temporary ASCII file, `ta_script.ssn`.

Next, the `nawk.st_cr1` procedure verifies that the simulation ended normally, that is, executed to the full requested simulation time, this being found in the original stability data file. If the time corresponding to the last recorded point does not equal this value, the case is rejected outright, and no other buses are examined. If the two are equal, the next criterion is called upon.

The `nawk.st_cr2` procedure applies the main voltage criterion, as specified in the ELISA script, and stores the results in the *case* summary file, `ob.`*topology_case*.

When all buses have been analysed individually, the result is analysed globally to establish case acceptability. For a given case to be acceptable, every monitored bus must have passed its individual criterion. Otherwise, it is rejected.

4.5 POST–CONTINGENCY PV–CURVE LIMIT SEARCH: INTERNAL PROCESS FLOW

The main advantage of the steady–state PV–curve limit search lies in the fact that only power–flow software is required to perform this analysis. Though hundreds of power–flow simulations may be needed to find a single PV–curve limit, these can be executed in a fraction of the time required for a single long–term stability simulation. This is significant in that a strong correlation has been found between voltage–stability limits obtained from long–term stability simulations and those resulting from steady–state post–contingency PV–curve analysis for at least one network [Mailhot, Gauthier & Vanier 1993]. The impact of such an approach is evident in the operations planning environment when many contingencies and degraded topologies must be considered. In fact, the mechanization of steady–state post–contingency PV–curve limit search processes within ELISA was instrumental in establishing the aforementioned correlation between long–term and steady–state post–contingency limits.

4.5.1 OVERVIEW OF PV–CURVE LIMIT SEARCH PROCESSES

The principles related to PV–curve limit determination outlined in 3.3.2.4 are well–known [IEEE 1991] will not be covered in great detail. The object of

this analysis is to find a transmission corridor's maximum capacity, usually in terms of the power P transmitted through the corridor and the voltage V at the corridor output (i.e. receiving end): This defines the P and V variables of the PV curve. To vary the voltage V associated with the receiving end, a fictitious generator (i.e. a PV bus) is created where real power output is set to zero.

The basic strategy for finding a PV curve is as follows: For every value of power dispatching P injected into the corridor, one varies the sending end voltage until the fictitious generator produces zero reactive power (i.e. Q = 0 MVAR): This condition coincides with that of the true network with no fictitious bus and consequently permits the receiving end voltage to be determined for the dispatched P. When this is repeated for many values of P, one can plot the PV curve and find the limit.

In ELISA, these concepts have been adapted to perform a PV–curve limit search in the following way:

1. In the ELISA script, the user identifies the process variables and criterion as outlined in 3.3.2.4.
2. ELISA performs a sweep of the fictitious bus voltage values for every value of power dispatching P identified within the search process.
3. If the fictitious bus reactive power generation Q changes sign in the course of the voltage sweep, this signifies an intersection with the network PV–curve (i.e. for the specified P and V variables). This also means that the power P is less than the maximum.
4. If the fictitious bus reactive power generation Q remains positive for all values of fictitious bus voltage, this means that the value of P is beyond the PV–curve maximum.
5. Having found at least one value of P less than the limit, and one other value greater than the limit, ELISA uses a binary search to find the limit within a predefined error of 100 MW.

4.5.2 OVERVIEW OF ELISA POST–CONTINGENCY PV–CURVE LIMIT SEARCH PROCESSES

Before one can perform a *post–contingency* PV–curve limit search, a number of additional issues must be addressed. The user must first prepare a base

case power flow where all EHV buses are defined as fictitious PV buses (see 2.4.7). As usual, the base case power flow should ideally reflect true, though conservative, network conditions. In particular, as described in [Mailhot, Gauthier & Vanier 1993], care must be exercised in:

1. modelling the voltage regulation characteristic of SVCs and SCs;
2. ensuring that off–load tap changers are not modelled as on–load transformer tap–changers;
3. ensuring that power stations (i.e. real PV buses) are controlling the voltage at a local bus (i.e. the power station's own), not that of a distant bus;
4. ensuring that power stations are modelled using correct reactive power limits (though SVCs and SCs are constrained to generate 100 MVARs in the pre–contingency power flow, see 2.4.7), and
5. changing the location of the swing bus, if necessary. On the one hand, one does not which to position the swing bus near the contingency as this will artificially support voltage in the neighbourhood of the contingency. On the other hand, a contingency will cause losses to increase in the affected corridor and one wishes to simulate the network's normal power–frequency reaction to increased dispatching requirements.

To find a post–contingency PV–curve limit, one must of course vary power dispatching. In ELISA, the power dispatching is changed by applying a new value for power on the pre–contingency network (i.e. with numerous fictitious PV buses): As already indicated in 2.4.7, this is not only to facilitate the convergence of the RP600 power flow software, it also corresponds to a fundamental assumption in network operations: The operator must ensure that a flat 1.0 p.u. voltage profile is maintained at all times, and perform corrective action if this is not the case. It must be consequently be assumed that the pre–contingency network has a 1.0 p.u. voltage profile, regardless of power transfer, which means that the simulated contingency is applied only after power dispatching has been changed. A consequence of this is that every voltage sweep is performed on a different PV curve (i.e. because passive shunt voltage support has changed to maintain a 1.0 p.u. voltage profile, thereby changing the topology). Nevertheless, there exists a pre–contingency network with 1.0 p.u. voltage profile which, when subjected to the contingency, will result in a post–contingency network

with no feasible load flow solution (i.e. if the fictitious bus were absent).

The contingency is stored in a standard RP600 power flow modification file. In order to properly simulate network conditions, this "contingency file" is used to perform a variety of functions [Mailhot, Gauthier & Vanier 1993]:

1. The contingency is identified.
2. The fictitious bus (i.e. where voltage is to be varied) is created.
3. All PV buses to be transformed into PQ buses are identified, and the reactive power generated or absorbed is modelled as an appropriate shunt impedance.
4. SVC and SC reactive power ranges are restored to their true values.
5. Shunt reactor switching is simulated.

As in the case of transient and long-term stability, the base case power flow data file is named using the `so.`*topology* format. The name of the contingency file is given the `co.`*topology* format where `co.` is the suffix and *topology* is the name of the particular topology for which a limit is to be found.

4.5.3 THE POST-CONTINGENCY PV-CURVE LIMIT SEARCH PROCESS: THE `csh.pc_sch` PROCEDURE

4.5.3.1 ENTERING THE LIMIT SEARCH PROCESS

Referring to Fig. 4-8 (a), the `csh.pc_sch` procedure begins by executing the load flow corresponding to the data supplied by the user in the `so.`*topology* file. This is the pre-contingency load flow: As mentioned before, this case is characterized by a flat voltage profile (usually at 1.0 p.u. voltage) due to the use of fictitious PV buses at every EHV bus (see 2.4.7). If the pre-contingency load flow does not converge, the process stops here as the data is assumed to be in error, and ELISA continues on to another topology.

If the pre-contingency load flow converges, the `csh.pc_lf` procedure reads the `co.`*topology* "contingency file" and applies these modifications to the pre-contingency load flow data: This executes a new power flow simulation where the main focus is 1) applying the contingency, 2) creating a fictitious bus where voltage is to be varied and 3) transforming all fictitious buses into PQ buses. As pointed out previously, shunt reactor switching may also be simulated,

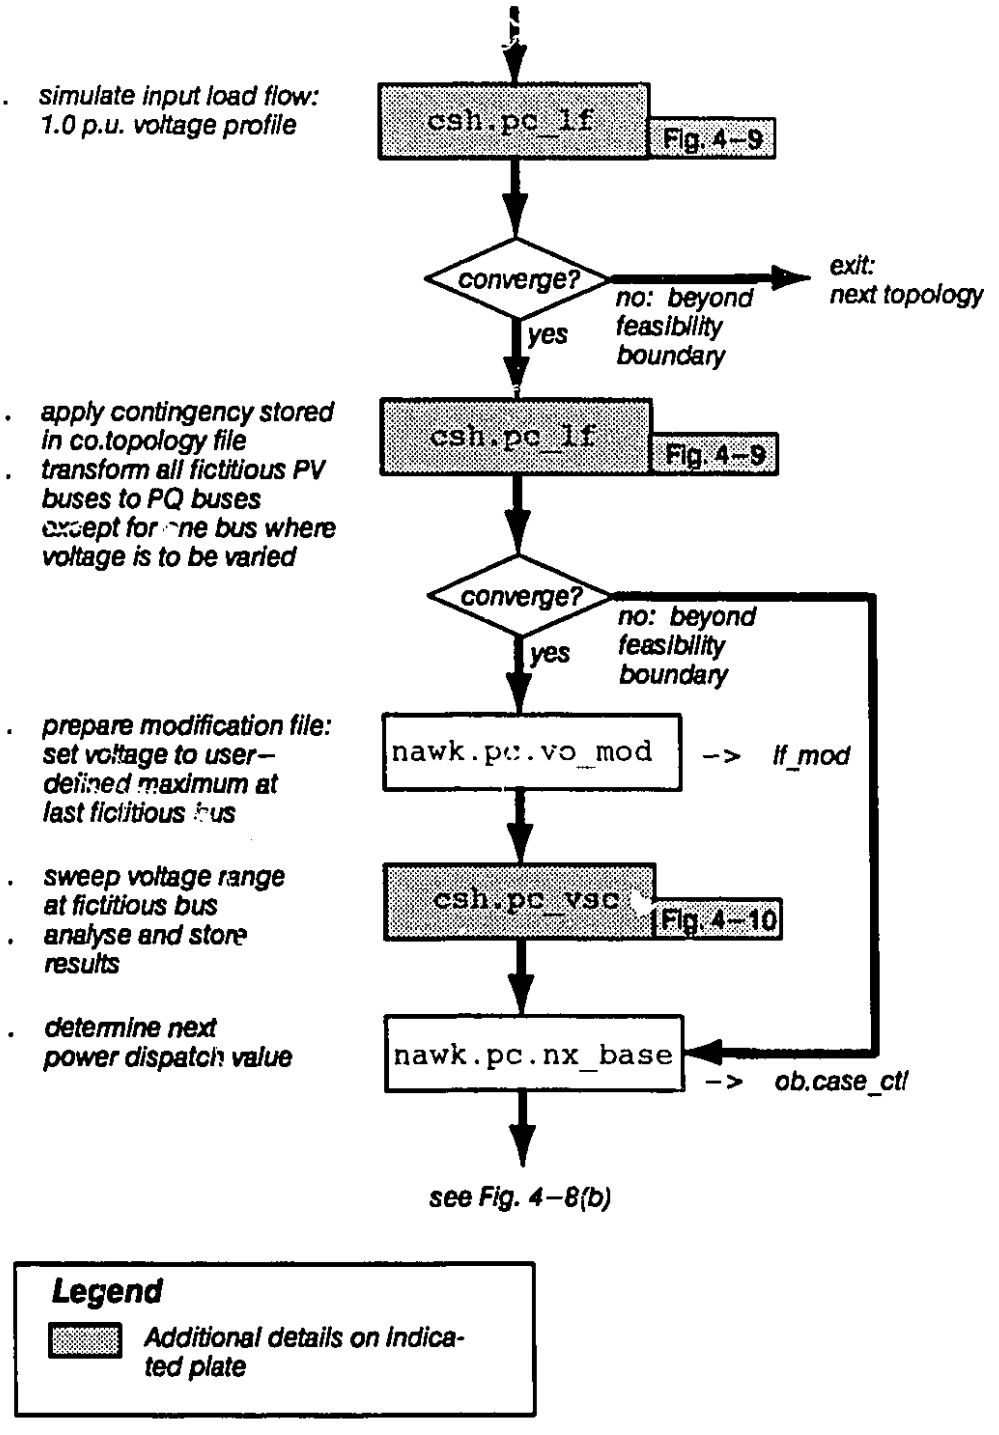

Fig. 4–8(a) ELISA process flow: The `csh.pc_sch` post–contingency PV curve limit search procedure.

and SVC and SC reactive power ranges are usually restored. For the sake of clarity, we will refer to this case as the initial contingency load flow.

If the initial contingency load flow does not converge, the load flow feasibility boundary has been crossed: There is therefore no need to sweep through voltages at the fictitious bus. The `nawk.pc.nx_base` procedure concludes that power dispatching must be reduced by the user–defined `init_inc`, and writes the new value for power dispatching in the `ob.case_ctl` file.

If the initial contingency load flow converges, the `nawk.pc.vo_mod` procedure creates a new modification file (i.e. `lf_mod`) whose objective is to set the voltage of the fictitious bus to the `maximum` user–defined value identified in the script. The `csh.pc_vsc` procedure then sweeps through the range of user–defined voltages at the fictitious bus: It begins with the `maximum` value, decrements through the voltages until the user–defined `minimum` has been reached, and then determines whether or not the PV–curve limit has been crossed. The `nawk.pc.nx_base` procedure then determines the search direction and identifies the next value for power dispatching based on the user–defined `init_inc`: This value is then written in the `ob.case_ctl` file.

4.5.3.2 THE MAIN LIMIT SEARCH PROCESS LOOPS

Referring to Fig. 4–8 (b), based on the previous analysis, appropriate case and file names are created. The `nawk.pc.lf_mod` procedure prepares a modification file (i.e. `lf_mod`) in order to change the power dispatching of the pre–contingency load flow case. The `csh.pc_lf` procedure then applies this modification file to the pre–contingency load flow data and performs a power flow simulation based on the modified data.

If, at this new value of power dispatching, the pre–contingency load flow does not converge, the load flow feasibility boundary has been crossed: There is therefore no need to sweep through voltages at the fictitious bus. The `nawk.pc.nx_case` procedure therefore concludes that power dispatching must be reduced by some value: This is determined by the binary search strategy, and the new value is written in the `ob.case_ctl` file.

If, at this new value of power dispatching, the pre–contingency load flow

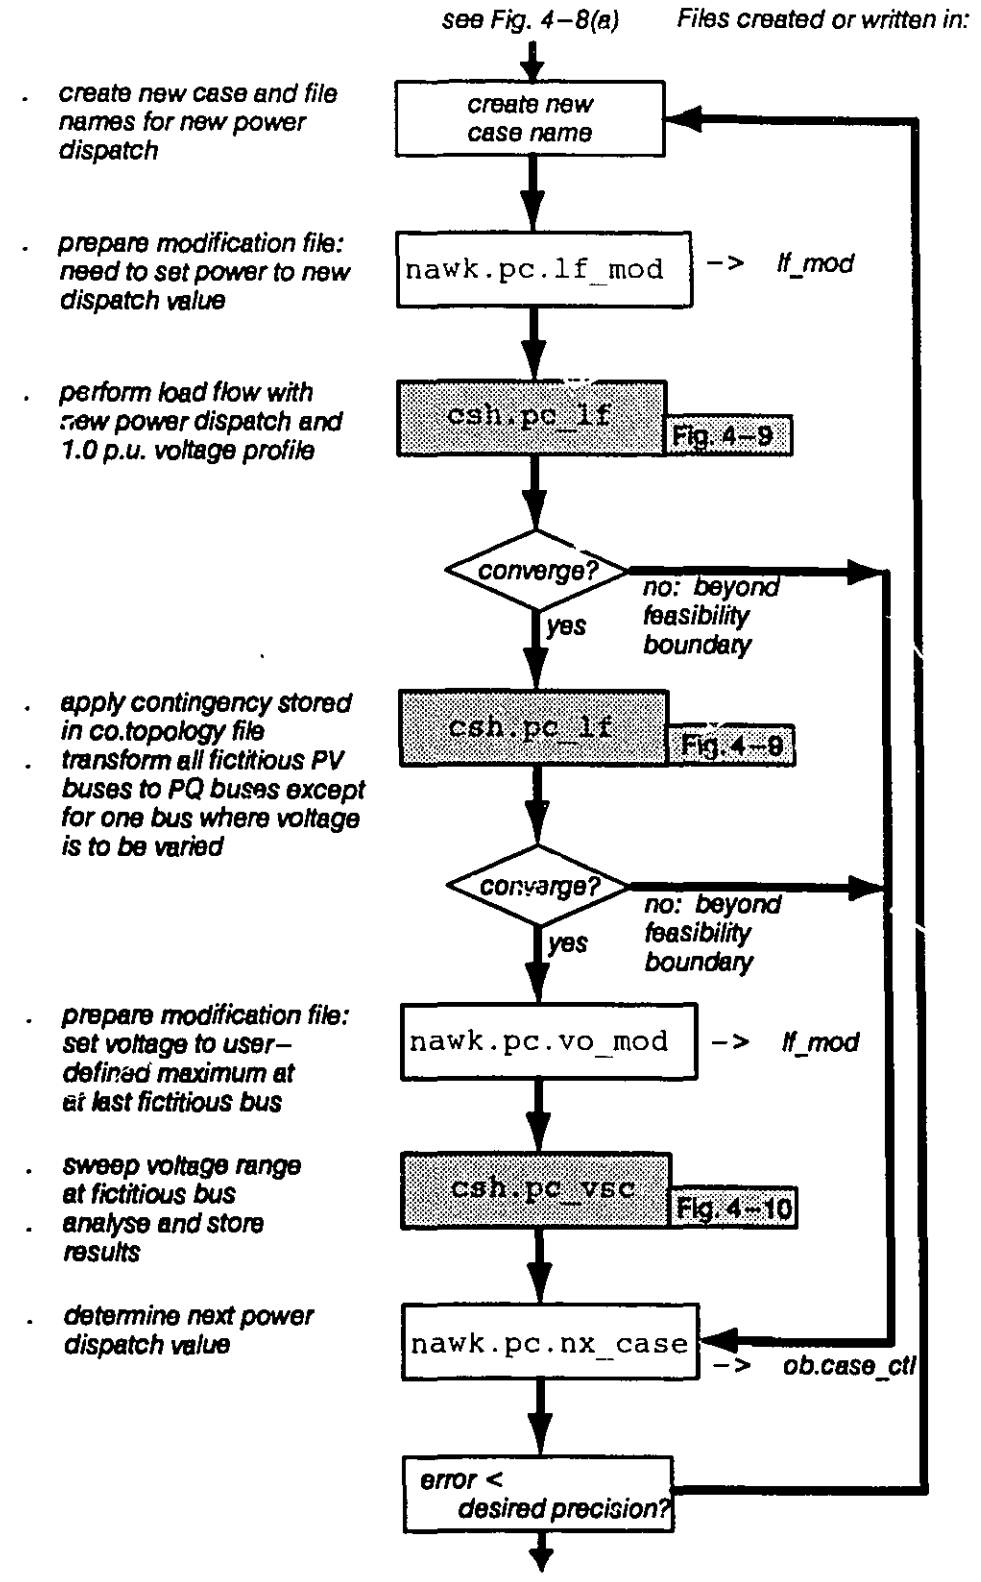

Fig. 4–8(b) ELISA process flow: The `csh.pc_sch` post–contingency PV curve limit search procedure.

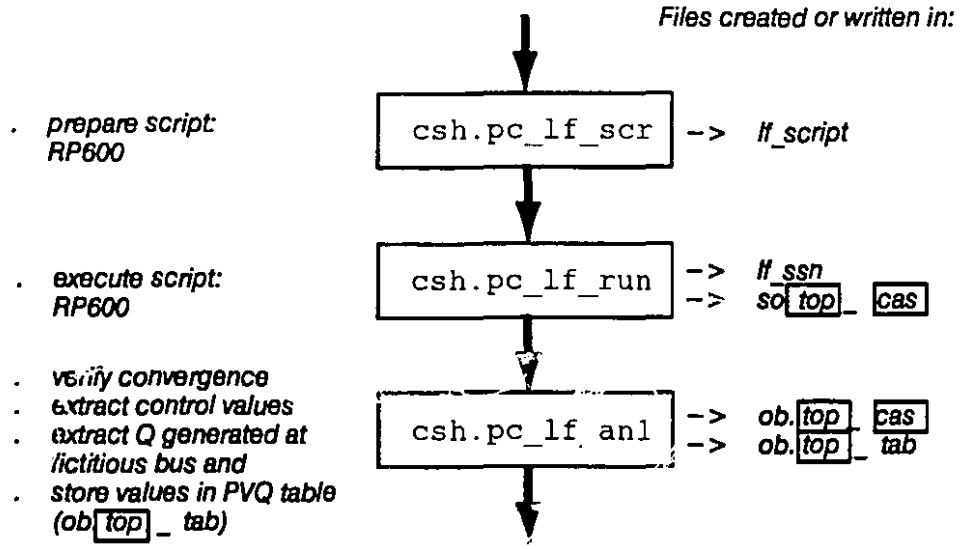

Fig. 4–9. ELISA process flow: The `csh.pc_lf` power flow (i.e. load flow) execution procedure.

converges, the `csh.pc_lf` procedure then reads the `co.`*topology* "contingency file" and applies this to the preceding case as before. For the sake of clarity, we will again refer to this case as the initial contingency load flow.

If the initial contingency load flow does not converge, the load flow feasibility boundary has been crossed: There is again no need to sweep voltages. The `nawk.pc.nx_base` procedure therefore concludes that power dispatching must be reduced by some value: This is determined by the binary search strategy, and the new value is written in the `ob.case_ctl` file.

If the initial contingency load flow converges, the `nawk.pc.vo_mod` procedure creates a new modification file (i.e. `lf_mod`) whose objective is to set the voltage of the remaining fictitious bus to the maximum user–defined value identified in the script. The `csh.pc_vsc` procedure once again sweeps through the range of user–defined voltages at the fictitious bus: It begins with the maximum value, decrements through the voltages until the user–defined minimum has been reached, and then determines whether or not the PV–curve limit has been crossed. The `nawk.pc.nx_case` procedure then determines

the next value for power dispatching based on the binary search strategy and the new value is written in the `ob.case_ctl` file. The process ends when the limit has been determined within 100 MW.

4.5.4 THE `csh.pc_lf` LOAD FLOW EXECUTION PROCEDURE

None of this will now be new to the attentive reader, the essentials having already been covered in 4.4.2. Referring to Fig. 4–9, the `csh.pc_lf_scr` procedure prepares the script required to execute the RP600 power flow software, and stores it in a file named `lf_script`. Once again, the script includes the name of the binary file in which results are to be stored (i.e. in the form `so.`*topology_case*).

The `csh.pc_lf_run` procedure is then called, executing the RP600 power flow according to the instructions contained in the script. The interactive power flow session, created in the course of the program's execution, is here again redirected and temporarily stored in the `lf_ssn` file. This file also contains the values of the control and adjusted variables requested in the RP600 script.

Finally, the `csh.pc_lf_anl` procedure analyses these results, first verifying for convergence of the power flow, then extracting `control` and `adjust` variable values in addition to the reactive power generated at the fictitious PV bus. These values are then stored in the `ob.`*topology_case* file. The power dispatching, voltage and reactive power generated at the fictitious bus are added to the result table found in the `ob.`*topology*`_tab` file.

4.5.5 THE `csh.pc_vsc` VOLTAGE SWEEP PROCEDURE

Referring to Fig. 4–10, the `csh.pc_lf` procedure performs a post–contingency load flow simulation, this time using the `lf_mod` modification file to redefine the voltage at the fictitious bus to the `maximum` user–defined value.

If the post–contingency load flow at `maximum` user–defined voltage does not converge, the load flow feasibility boundary has been crossed and we exit this procedure: The binary search will set the next power dispatching to a lower value.

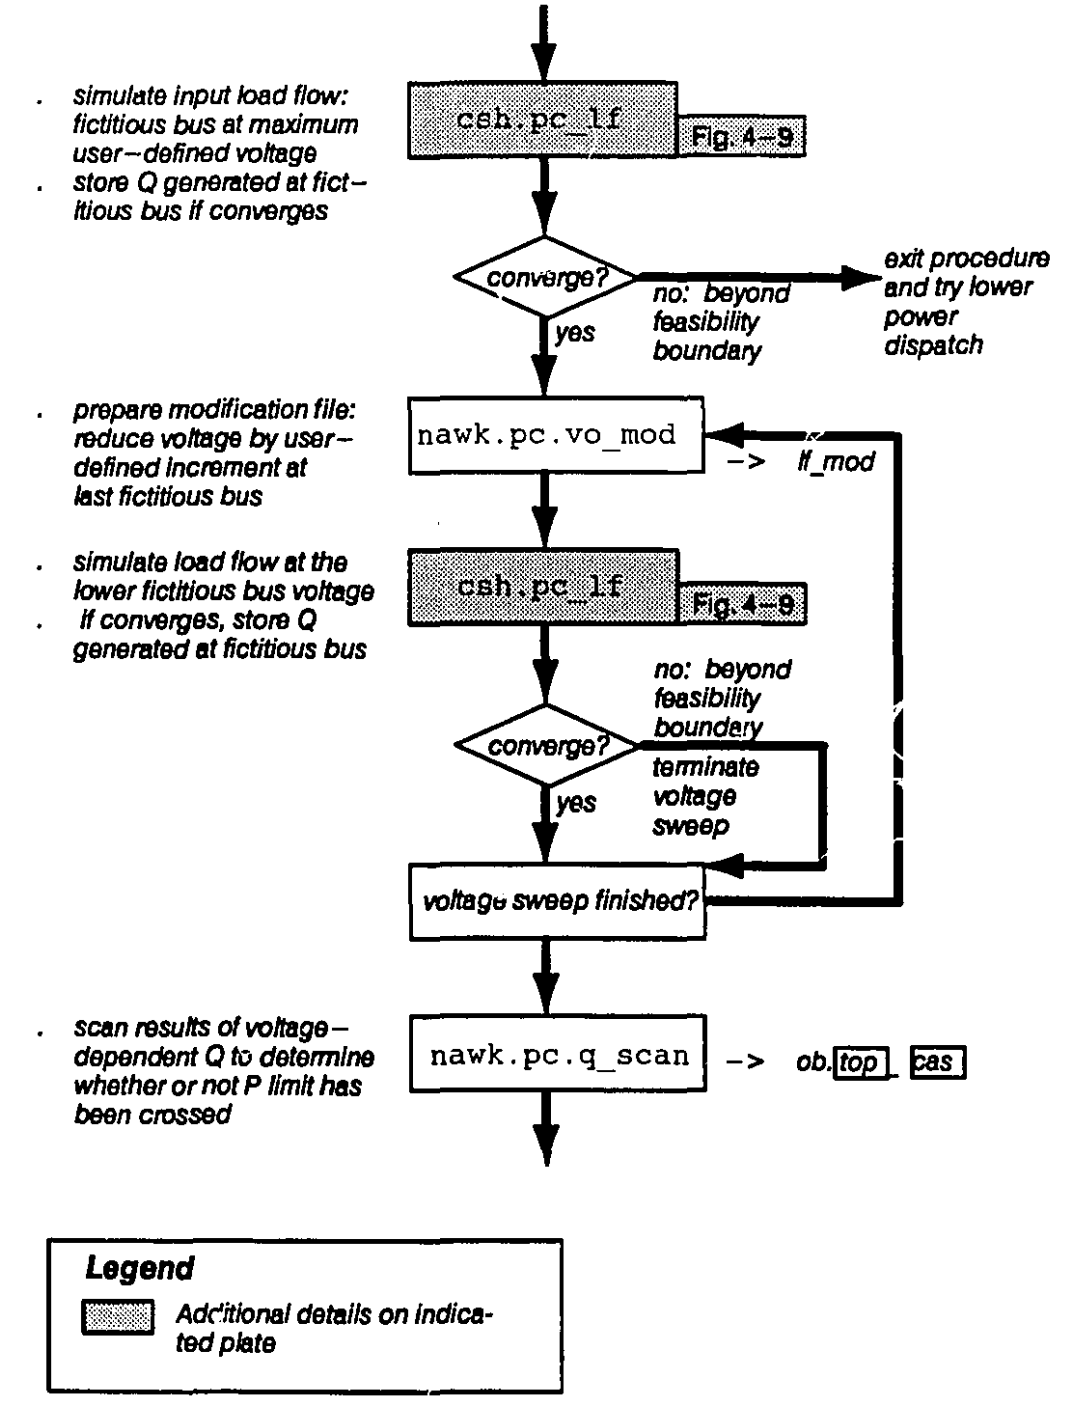

Fig. 4–10 ELISA process flow: The `csh.pc_vsc` voltage sweep procedure which determines relative position with respect to the PV curve limit.

If the post–contingency load flow at maximum user–defined voltage converges, the nawk.pc.vo_mod procedure prepares a modification file which reduces the voltage at the fictitious bus by the user–defined increment. The csh.pc_lf procedure then simulates the load flow at this lower value of fictitious bus voltage.

If the lower–voltage post–contingency load flow simulation does not converge, the feasibility boundary has been crossed and the voltage sweep is terminated. If the lower–voltage post–contingency load flow simulation converges, fictitious bus voltage is tested to determine whether the full range has been scanned. If not, the process loops back to the nawk.pc.vo_mod procedure and continues. If the full range has indeed been scanned, the nawk.pc.q_scan procedure then examines the voltage–dependent reactive power output for the range of voltages which has just been scanned and determines whether or not the limit has been crossed.

4.6 EXAMPLES OF SCRIPTS SUPPORTED BY ELISA

The ELISA framework prototype has been in use for over a year and a half, in various production versions, by operations planning personnel at Hydro–Québec. ELISA provides considerable flexibility in defining various types of limit searches and determining transfer limits. In addition to this, the mechanization of processes in dynamic security analysis provides a formidable leverage on individual performance. Extremely large and ambitious studies can now envisaged and, more importantly, realized, in a fraction of the time previously required. The present section, and the next, provide some feel for these aspects.

A final remark. The examples presented below systematically use the degraded Hydro–Québec network described in section 2.3 of Chapter 2. The use of different topology names in the example scripts refers only to different contingency scenarios defined in the respective stability data files.

4.6.1 A TYPICAL GENERATION–LOAD LIMIT SEARCH

The results shown in Fig. 2–2 for the example presented in Chapter 2 were in fact generated by ELISA using the script shown in Fig. 4–11. This corresponds to a typical generation–load limit search. The main feature here is that ELISA uses the simulate command to build the contingency directly. The initial increment, init_inc, causes power transfer to increase by 800 MW after

```
for topology
     dtop2
{ find limit under transient stability with  precision
                  st600                          100
  { simulate for fault of duration
                              6
    for bus with  fault_volt and  loss_of_line
        731         .739            731 717 1
  { control generation of zone  weight with init_inc
                           59     1            800
  and adjust load of network with adjust_factor
                                          0.9
  and accept as a function of voltage of  bus min max after
                                          702 .95 1.1  500
  }
  }
  and trace with scale of voltage min max
                                  0.9 1.0
}
end
```

Fig. 4–11. Script used to perform the limit search example of Fig. 2–2.

simulating the base case, and the search strategy thereafter switches to the binary search. The transfer limit is higher than the base case.

4.6.2 A GENERATION–GENERATION LIMIT SEARCH WITH SLG FAULT IMPEDANCE DETERMINATION

Fig. 4–12 illustrates two features not yet shown. First, the script requests that the SLG fault impedance be calculated before attempting to find the `limit` for topology `hq2` (whose power flow and stability data are found in the files respectively named `ma.hq2` and `so.hq2`). This script also requests that a generation–generation limit search between the LG 2 (i.e. bus 49) and LG 3 (i.e. bus 64) power stations of the Hydro–Québec network (Fig. 2–1) be performed where the objective is to increase the power flow on the line separating the two for a fault at LG 3: In this case, generation is `controled` (i.e. increased) at LG 3 and `adjusted` (i.e. decreased) at LG 2. We see that the `adjust` quantity performs negative adjustments with respect to the controlled quantity because of the *negative* `adjust_factor`.

The initial increment chosen in this case is 200 MW. From the results shown in Fig. 4–13, we see that the limit is found to be at a higher value of power trans-

```
for topology
     dtop2b
{ determine SLG fault impedance
and find limit under  transient stability with  precision
                        st600                          100
  { control generation of  bus weight with init inc
                          64    1            200
  and adjust generation of  bus weight with adjust factor
                           49    1              -1
  and accept as a function of voltage of  bus min max after
                                          702 .95 1.1 500
  }
  and trace with scale of voltage min max
                                    0.9 1.1
}
end
```

Fig. 4–12. Example of script employing a generation–generation limit search strategy. SLG fault impedance determination is also requested.

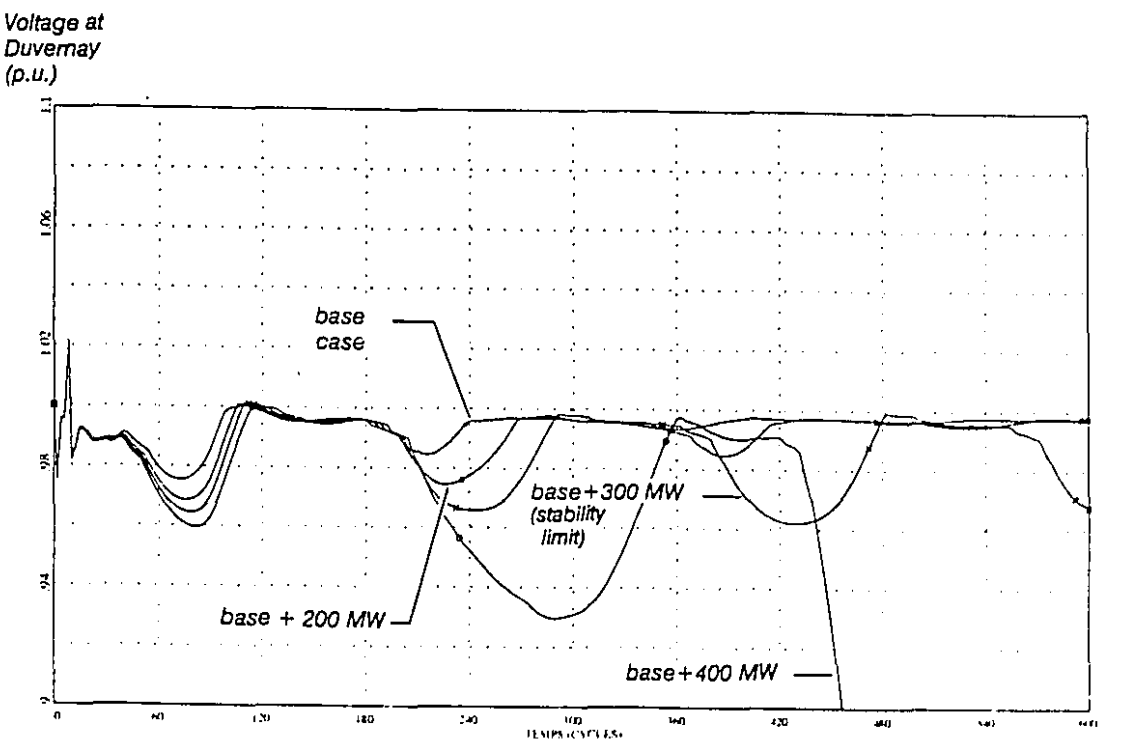

Fig. 4–13. Graphical output generated by ELISA for the above script (using the STPP post–processor software for ST600 stability output). The voltage at Duvernay (bus 702) is plotted as a function of time for different values of power transfer.

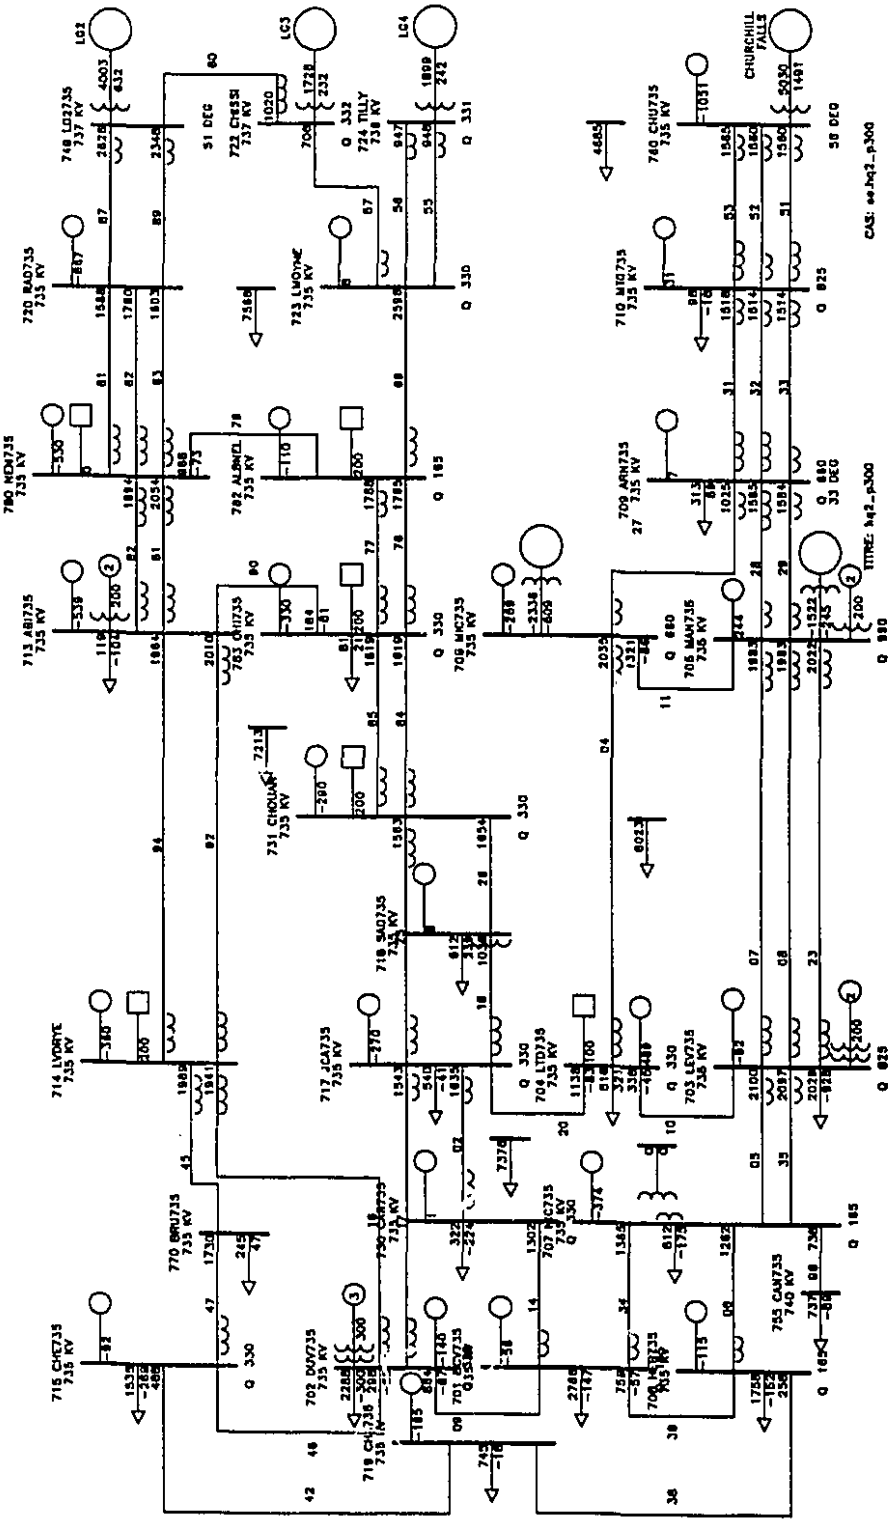

Fig. 4–14. Single line diagram of power flow of transfer limit case (script of Fig. 4–12). Such a diagram is generated by ELISA using the RP600 power flow software immediately following a transient stability limit search.

fer than the base case. In this case, the limit search process first attempted the base case (acceptable), base+200 MW (acceptable) and base+400 MW cases: Only the latter was found to be *unacceptable*. Initiating now the binary search, the base+300 MW case was attempted and found to be acceptable. At this point, the distance between unacceptable and acceptable cases being equal to the desired `precision` (i.e. 100 MW), the search was interrupted. Fig. 4−14 shows the single line diagram of the power flows for the limit (i.e. base+300 MW), automatically generated by ELISA using the RP600 software.

4.6.3 A LOAD−GENERATION NETWORK LIMIT SEARCH

Fig. 4−15 shows a script where the main feature is that load is `controled` and generation is `adjusted` on the entire network. Because `generation` is `adjusted`, the `adjust_factor` is greater than 1. In addition to this, the generation and load changes are made on the entire network. An almost identical search could have been performed by `controling generation` and `adjusting load` with an `adjust_factor` of `0.9`.

Examining the resulting waveforms of the script−requested fault scenario (i.e. SLG fault at Radisson − bus 720 − with loss of line to Nemiscau) at the limiting bus (i.e. in this case, this is at La Vérendrye, bus 714), Fig. 4−16 shows that the limit is lower than the base case, which proves that ELISA's modified binary search is effective in all directions.

4.6.4 A POST−CONTINGENCY PV−CURVE LIMIT SEARCH

As a final example, Fig. 4−17 shows a script describing a typical post−contingency PV−curve limit search. In the script, three `topologys` are considered: `lvd1_pc`, `lvd2_pc` and `lvd3_pc`. Additionally, generation is controled in `zone 59` and `load` is `adjusted` on the entire network. We see that the criterion requires accepting results on the basis of 0 MVAR at fictitious `bus 717`, and the range of `voltages` to be considered at every value of power dispatching is between `690` and `735` kV. The voltage sweep is performed in `increments` of `3` kV. The contingency associated with each topology is provided by the user in the `co.lvd1_pc`, `co.lvd2_pc` and `co.lvd3_pc` files.

Fig. 4−18 shows typical output generated by ELISA, illustrating the basic principles: These results were recorded in the `ob.lvd3_pc_tab` file. The first column shows the relative position of the power dispatching with respect to the

```
for topology
      hq1
{ find limit under transient stability with  precision
                  st600                          100
  { simulate for fault of duration
                                6
    for bus with  fault volt  and loss of line
        720          .625          720 780 1
  { control load of  network with  init inc
                                      800
  and adjust generation of  network with  adjust_factor
                                            1.1
  and accept as a function of voltage of  bus min max after
                                          702  .95 1.1 500
                                          703  .95 1.1 500
                                          714  .95 1.1 500
  }
  }
  and trace with scale of voltage min max
                                    0.9 1.1
}
end
```

Fig. 4–15. Example of script with load–generation limit search strategy.

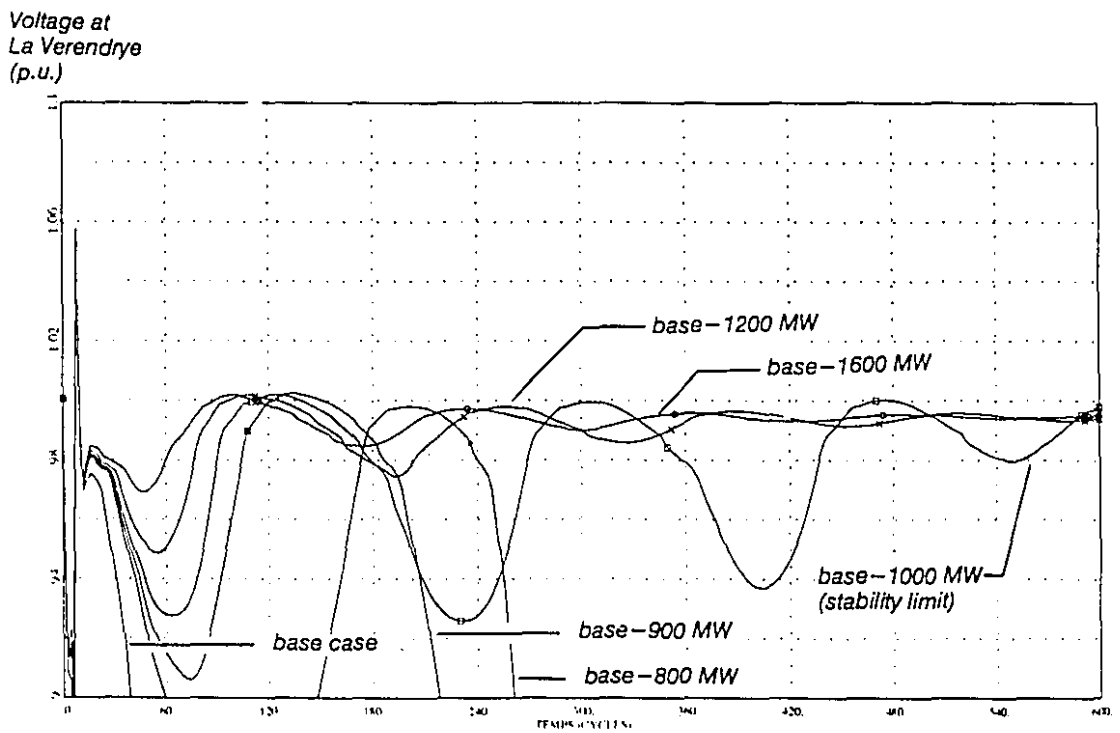

Fig. 4–16. Plot of voltage at La Vérendrye (bus 714) generated by ELISA.

```
for topology
      lvd1_pc
      lvd2_pc
      lvd3_pc
{ find limit under steady state with precision
                  rp600              100
      { control generation of  zone with init inc
                                 59          400
      and adjust load of  network
      and accept  as a function of mvar of  bus min
                                                717  0
                  for voltage min max inc
                                690 735  3
      }
  }
end
```

Fig. 4–17. Example of script for post–contingency PV–curve limit search.

base case. The second column indicates whether or not the particular case converged. The third column gives the voltage at the fictitious bus and the fourth gives the reactive power generated at the fictitious bus.

The table is constructed chronologically, clearly showing the evolution of the binary search strategy. In this limit search example, the base case (i.e. 0 MW) shows only positive reactive power values, indicating that the case is beyond the limit. The base case – 100 MW shows two negative reactive power values just before arriving at the feasibility limit: This is consequently taken to be the limit.

4.7 EXAMPLES OF TYPICAL SENSITIVITY STUDIES

4.7.1 TRANFER LIMIT DETERMINATION AS A FUNCTION OF GENERATION REJECTION

To illustrate the impact of mechanizing dynamic security analysis processes using frameworks, a sensitivity analysis was performed on rejecting increasing numbers of generating units at the LG 2 power–station for an SLG fault at the Abitibi 735 kV switching station with loss of line to La Vérendrye. The SLG fault is applied at 0 cycles, with clearing of the fault occurring at 6 cycles. Generation rejection is initiated at 16 cycles, and automatic reclosure on the faulted line section is performed at 85 cycles, resulting in circuit–breaker operation and

Topologie: lvd3_pc

Position (MW)	Converg.	Tension (kV)	MVAR
0	oui	735.0	637.5
0	oui	732.0	554.4
0	oui	729.0	455.5
0	oui	726.0	361.1
0	oui	723.0	306.1
0	oui	720.0	209.5
0	oui	717.0	131.1
0	oui	714.0	94.8
0	non		
-400	oui	735.0	232.2
-400	oui	732.0	144.2
-400	oui	729.0	58.4
-400	oui	726.0	-56.6
-400	oui	723.0	-138.6
-400	oui	720.0	-213.8
-400	oui	717.0	-282.4
-400	oui	714.0	-381.9
-400	oui	711.0	-459.7
-400	oui	708.0	-532.2
-400	oui	705.0	-511.3
-400	oui	702.0	-547.6
-400	oui	699.0	-463.4
-400	oui	696.0	-330.6
-400	oui	693.0	229.7
-400	non		
-200	oui	735.0	410.9
-200	oui	732.0	325.4
-200	oui	729.0	225.5
-200	oui	726.0	134.3
-200	oui	723.0	46.7
-200	oui	720.0	-31.9
-200	oui	717.0	-96.5
-200	oui	714.0	-198.9
-200	oui	711.0	-276.1
-200	oui	708.0	-281.1
-200	oui	705.0	-146.6
-200	oui	702.0	101.2
-200	non		
-100	oui	735.0	534.8
-100	oui	732.0	449.8
-100	oui	729.0	357.9
-100	oui	726.0	269.7
-100	oui	723.0	191.1
-100	oui	720.0	127.9
-100	oui	717.0	0.3
-100	oui	714.0	-88.2
-100	oui	711.0	-120.8
-100	oui	708.0	37.8
-100	non		

Fig. 4–18. Example of the type of results generated by ELISA for a post–contingency PV–curve limit search. In this case, the results were stored in the `ob.lvd3_pc_tab` file.

the permanent loss of a line 6 cycles later.

The simulations were run for 300 60−Hz cycles. The script used to initiate the study is presented in Fig. 4−19. As usual, each of the four topologies is identical except for the number of generating units to be rejected in the stability data file of each one. For the purposes of this study, the minimum acceptable voltage was set to 0.94 p.u. after 150 cycles at the faulted station (i.e. at Abitibi, bus 713).

Figure 4−20 shows typical Abitibi voltage waveform generated by the study: Here we see the two−unit rejection scenario at LG 2. Each waveform corresponds to a different value of pre−contingency power transfer in the James Bay corridor. In this instance, only the base case meets the stated minimum voltage criterion: We take this to be the acceptability limit.

The results of this study are shown in Fig. 4−21. The first unit rejected improves power transfer by 500 MW, and each succeeding unit increases power transfer by 300 MW. These results were generated after executing 20 stability cases in less than five hours on a SUN Sparcstation 2, each transfer limit search having required 5 stability simulations. This study would have required, on average, approximately two to three days for a typical operations planner.

4.7.2 SECURITY LIMIT DETERMINATION AS A FUNCTION OF LOAD REPRESENTATION

A sensitivity study on fault location was performed to find the worst contingency and the resulting security limit on the James Bay transmission system (Fig. 2−1). Network load representation was modified by changing, in the ST600 stability input data file for each topology, the value of the voltage exponent, *alphap*, between 1 and 1.5 in the following equation for real power at all constant−power load buses:

$$P = P_0 \times V^{alphap} \qquad (4-1)$$

An *alphap* of 1 is typically used to simulate a summer load on the Hydro−Québec network, and 1.5 yields a good representation of winter peak load due to extensive use of domestic electric heating. The variable P_0 is the value of the load provided by the pre−contingency power flow solution at each bus.

```
for topology
      r01g
      r11g
      r21g
      r31g
{ find limit under transient stability with precision
                    st600                          100
  { control generation of zone  weight with init_inc
                          59      1            800
  and adjust load of  network with  adjust_factor
                                          0.9
  and accept as a function of voltage of  bus min max after
                                           713  .94 1.1  150
  }
  and trace with scale of voltage min max
                                      0.7 1.1
}
end
```

Fig. 4–19. Script used to perform a sensitivity study on the number of units rejected at LG 2 in response to a complex contingency.

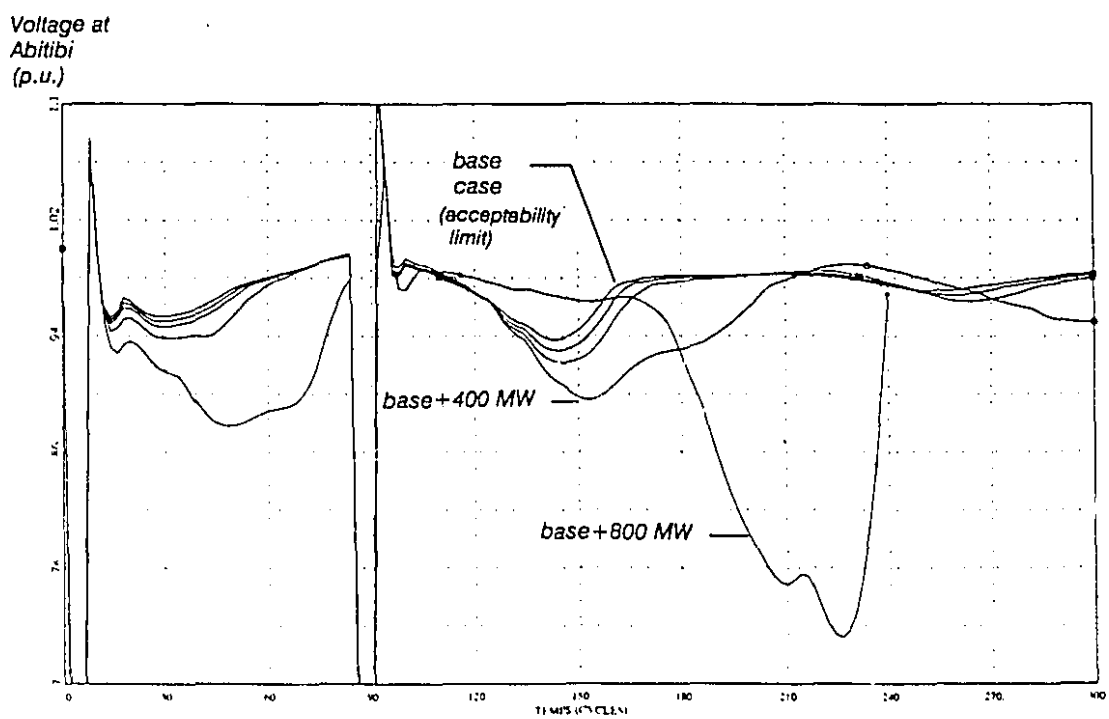

Fig. 4–20. Voltage at Abitibi for 2–unit rejection scenario at LG 2.

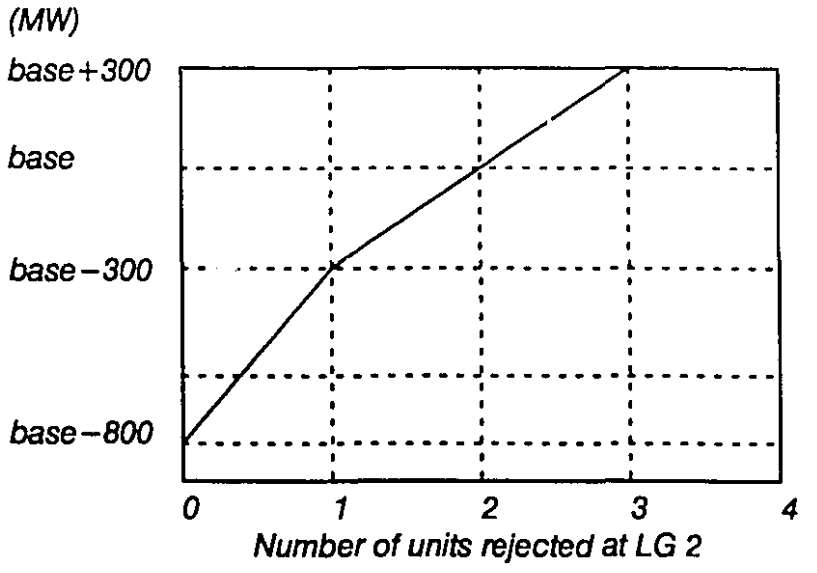

Fig. 4−21. James Bay 735 kV transmission capacity as a function of the number of units rejected at LG 2.

The script of Fig. 4−22 shows how ELISA was asked to find a power transfer limit at each station on the James Bay transmission corridor. This script was generated for three identical topologies whose only difference was the value of *alphap*. Zone 59 includes the LG 2, LG 3 and LG 4 power stations feeding the faulted corridor and bus 702 represents Duvernay switching station.

In a single twelve−hour evening period, the ELISA search processes executed 140 stability cases, running under three separate windows with remote logins on three SUN Sparcstation 2*s*. Each script ran under a different workstation and generated, on average, more than 46 stability cases. *A typical operations planner could have taken anywhere between two and six months to do the same work.*

Figure 4−23 shows the variation of power transfer limits on the James Bay system as a function of fault location. In particular, one sees that the network is weaker at Némiscau due to the absence of two SVCs there. One also sees that the worst contingency shifts in location according to winter or summer loading, and that the value of the security limit also changes. For instance, the worst contingency for winter loading is given by the fault at Radisson at *base − 400 MW*

```
for topology
      dtop2
{ find limit under transient stability with  precision
                    st600                        100
  { simulate for fault of duration
                                6
    for bus  with fault_volt  and loss_of_line
        720          .625            720 780 1
        723          .707            723 782 1
        780          .655            780 713 1
        782          .661            782 783 1
        713          .649            713 714 1
        783          .653            783 783 1
        731          .739            731 717 1
        714          .658            714 770 1
    { control generation of  zone  weight with init_inc
                              59      1            800
    and adjust load of network with adjust_factor
                                              0.9
    and accept as a function of voltage of  bus min max after
                                              702 .95 1.1 500
    }
  }
  and trace with scale of voltage min max
                                    0.9 1.1
}
end
```

Fig. 4–22. Script for performing large–scale sensitivity study of transfer and security limits as a function of seasonal load representation.

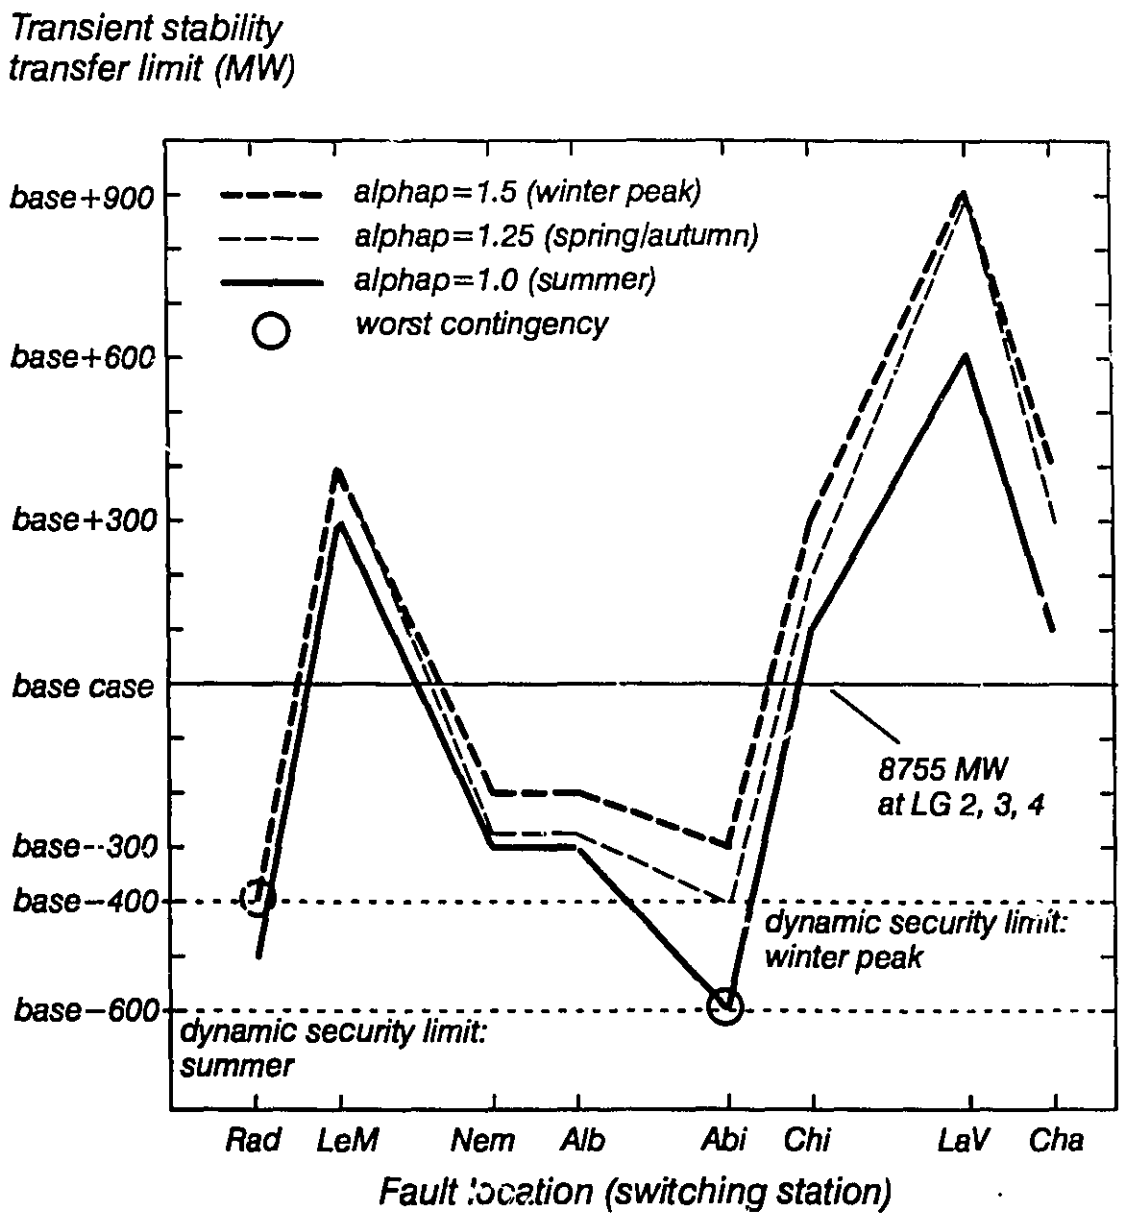

Fig. 4−23. Distribution of transient stability transfer limits on James Bay network as a function of fault location and *alphap* (SLG fault with loss of line on four−line topology). Note change in worst contingency and security limit with season. Note also the effect of two missing SVCs at Némiscau.

whereas the one for summer loading is given by the fault at Abitibi at *base – 600 MW*. In addition, one sees that the higher value of *alphap*, corresponding to winter loading, has a stabilizing effect on the network. This effect also increases as the fault draws closer to the load.

A final point. Power transfer limits corresponding to an *alphap* of 1.25 seem to coincide, at times, with limits for either one of the other two *alphap*. This is because all values were found within 100 MW. A smaller tolerance would have permitted greater discrimination.

4.8 SUMMARY

The present chapter provides a overview of the ELISA framework prototype. The internal ELISA software process structure has been described and a discussion on the choice of development tools has also been provided. Essentially, ELISA was developed directly in the language of the UNIX C–shell operating system because operating systems are at the top of the software control pyramid: Any software that can be executed under a particular operating system can be seen as *de–facto* integrated within this operating system. Programmable operating systems can be seen as a means for creating ever higher–level software from elementary building blocks which are neither elementary nor meant to be used as "building blocks".

From the point of view of dynamic security analysis, ELISA mechanizes routines traditionally carried out by experts, greatly accelerating the realization of complex processes. The user describes dynamic security analysis processes by means of the object–oriented language formally described in the previous chapter. The semantic net not only provides a rigourous definition of the ELISA language, the shaded portions clearly identify the aspects supported by the ELISA prototype and provides a roadmap to future enhancements.

To illustrate the flexibility and performance of the ELISA prototype, a number of example scripts are presented, including those of very large sensitivity studies. It is seen that the prototype permits the execution of large–scale studies in a fraction of the time formerly required, and enables operations and system planners to work at a more strategic level.

PART 2

THE ELISA FRAMEWORK AS A RESEARCH TOOL

CHAPTER 5

ON ESTIMATING POWER SYSTEM STABILITY LIMITS

5.1 INTRODUCTION

Part 1 of this thesis has focused on the mechanization of dynamic security analysis in operations planning. In particular, the ELISA framework prototype was described. ELISA was found to give extremely effective performance in mechanizing transient and long–term stability transfer limit determination processes using existing commercial simulation software packages even though a simple binary search strategy was employed.

However, is it possible to improve on binary search performance? As a prerequisite to this, a deeper understanding of network behaviour is required, particularly as it approaches instability. Specifically, one must be capable of quantifying the relative severity of a contingency for different network operating points and that of different contingencies.

Of course, the search for quantified measures of contingency severity is as much an art as a science. For each promising severity model, many simulations must be performed and the appropriate quantity extracted before trends are found and conclusions can be made. In order to concentrate on the high–level goal of finding such measures, it is important to be freed of the numerous, extremely time–consuming tasks which represent the principal obstacle to such an effort. This is precisely the area where the contribution of the ELISA framework proved to be decisive: Though conceived as a practical tool for operations planners, ELISA was modified only slightly to compute various target indices after having performed limit–search stability simulations. Individual severity indices were subsequently studied and plotted when appropriate. For a time, ELISA became a laboratory test bed in the search for an appropriate severity index.

As a result of this work, the present chapter shows that the signal energy of a network's transient response acts as a barometer which measures the relative severity of any normal contingency with respect to power generation or transfer:

For a given contingency, as power flow is increased and the network approaches instability, signal energy increases smoothly and predictably towards an asymptote which defines the network's stability limit. Additionally, the latter permits us to compare the severity of different contingencies. We shall see that this behaviour can be explained in terms of the effect of increasing power on the damping component of dominant poles. In particular, a simple function can be derived which estimates network stability limits with surprising accuracy from two or three simulations, provided that at least one of these is within 3 to 5% of the limit, depending on the type of contingency. These results hold notwithstanding the presence of many active, non−linear voltage−support elements (i.e. generators, synchronous condensers, SVCs, static excitation systems, etc.) in the network.

Before we address these issues, let us momentarily disgress and consider a matter which will be of great value in the coming analysis, that of the network impulse response.

5.2 Approximating the Network Impulse Response

One of the most meaningful approaches in characterizing any system is to begin with a study of its transient response. Let us once again consider the degraded Hydro−Québec network of chapter 2 and the stability limit search illustrated in Fig. 5−1 for a 6−cycle, 3−phase fault at Abitibi with no loss of a line. This contingency corresponds for example to a station bus fault followed by the loss of a short bus section within the station upon fault clearing. As we are now used to seeing, each 600−cycle waveform shows the time−dependent voltage behaviour $v(t, P)$ for different initial values of the steady−state generation P at James Bay (i.e. the combined LG2, LG3 and LG4 power stations); for the purposes of this search, the load was adjusted on the entire network. Figure 5−1 shows the waveforms for the search monitored at two locations, Duvernay and Abitibi. In all cases, pre−fault voltage is 1.0 p.u. and, for all stable cases, steady−state voltage eventually returns to 1.0 p.u. because the contingency includes essentially no post−fault change in topology. The stability limit is found within 25 MW.

Let us now suppose that a network is a black box with a single input and a single output (SISO): The user chooses where the input is applied and where

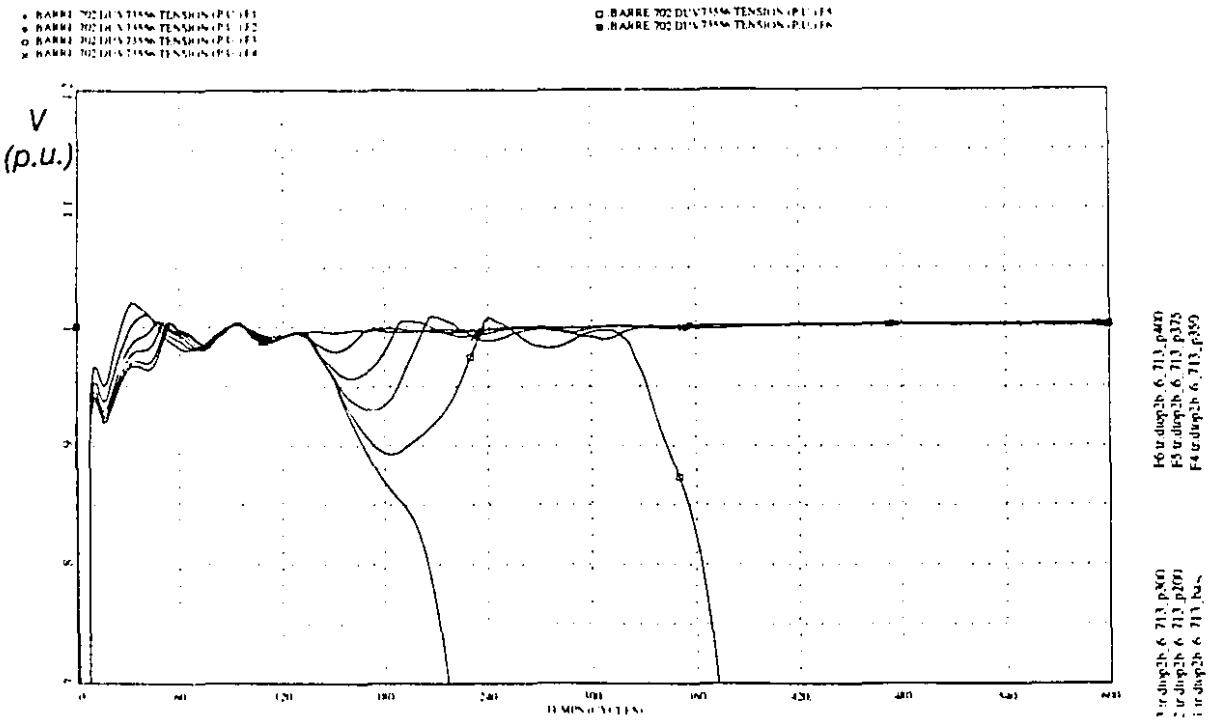

(a) Voltage monitored at Duvernay

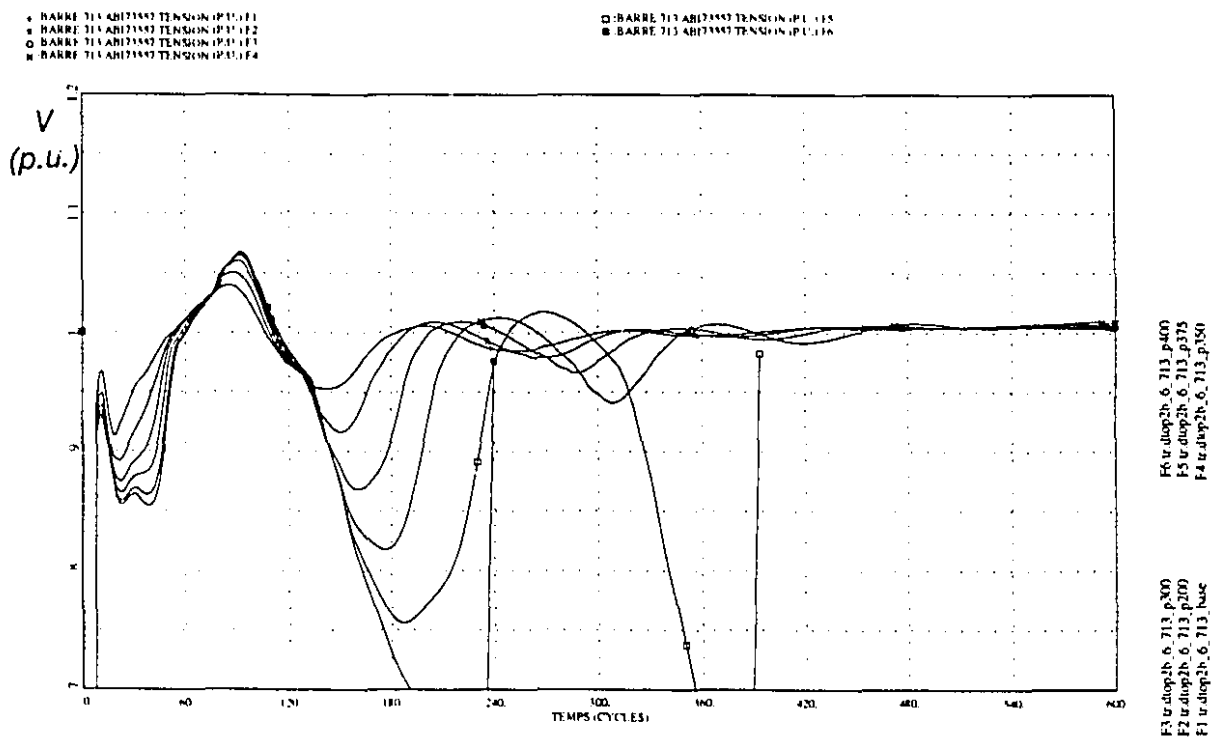

(b) Voltage monitored at Abitibi

Fig. 5−1. Transient stability transfer limit search performed by ELISA for 6−cycle, 3−phase fault at Abitibi with no loss of line.

the output is monitored. Let us also suppose that the black box has a knob which permits us to change some parameter, for instance, generation, which we will denote as P. If it were reasonable to assume that a fault, at a user–defined location, can be used to approximate an impulse at the input of this black box, then it might also be reasonable to assume that the quantity monitored at the output, for instance the voltage at some bus, approximates the impulse response h of this system (i.e. as seen from the user–defined monitoring location). Pursuing this hypothesis further, each of the waveforms of Fig. 5–1 could then be considered an approximation of the network impulse response h for different values of the parameter P. Let us now carefully examine this hypothesis.

5.2.1 REST STATE AND TRANSIENT RESPONSE

To obtain the impulse response of a linear system by simulating some appropriate impulse, the system must initially be at rest [Frederick & Carlson 1971]. In the case of each of the $v(t, P)$ waveforms of Fig. 5–1, let us assume the 1.0 per unit pre– and post–contingency steady state to be an appropriate rest state. The transient response $r(t, P)$ with respect to this rest state is therefore given by:

$$r(t, P) = v(t, P) - 1.0 \qquad (5-1)$$

5.2.2 APPROXIMATING THE IMPULSE AND ITS RESPONSE

The unit impulse is a singularity function often approximated by a rectangular pulse of amplitude $1/\varepsilon$ and width ε (i.e. such that the area is unity). As indicated by Frederick and Carlson, if ε is sufficiently small, that is, if the pulse is "short enough", the result will very nearly equal a true unit impulse, and one can experimentally – or by simulation – measure a system's impulse response. It now remains to establish the conditions which must be fulfilled by such a pulse, and the system on which it is used, for this to apply.

Consider the frequency domain. If some arbitrary pulse $\delta_a(t)$ is applied at the input of a linear system, its transfer function $H(\omega, P)$ is given as follows (see Fig. 5–2):

$$H(\omega, P) = R(\omega, P) \,/\, \mathcal{F}[\delta_a(t)] \qquad (5-2)$$

where $R(\omega, P)$ is the Fourier transform of the transient response $r(t, P)$ observed at the output monitoring location, and $\mathcal{F}[\delta_a(t)]$ is the Fourier transform of the

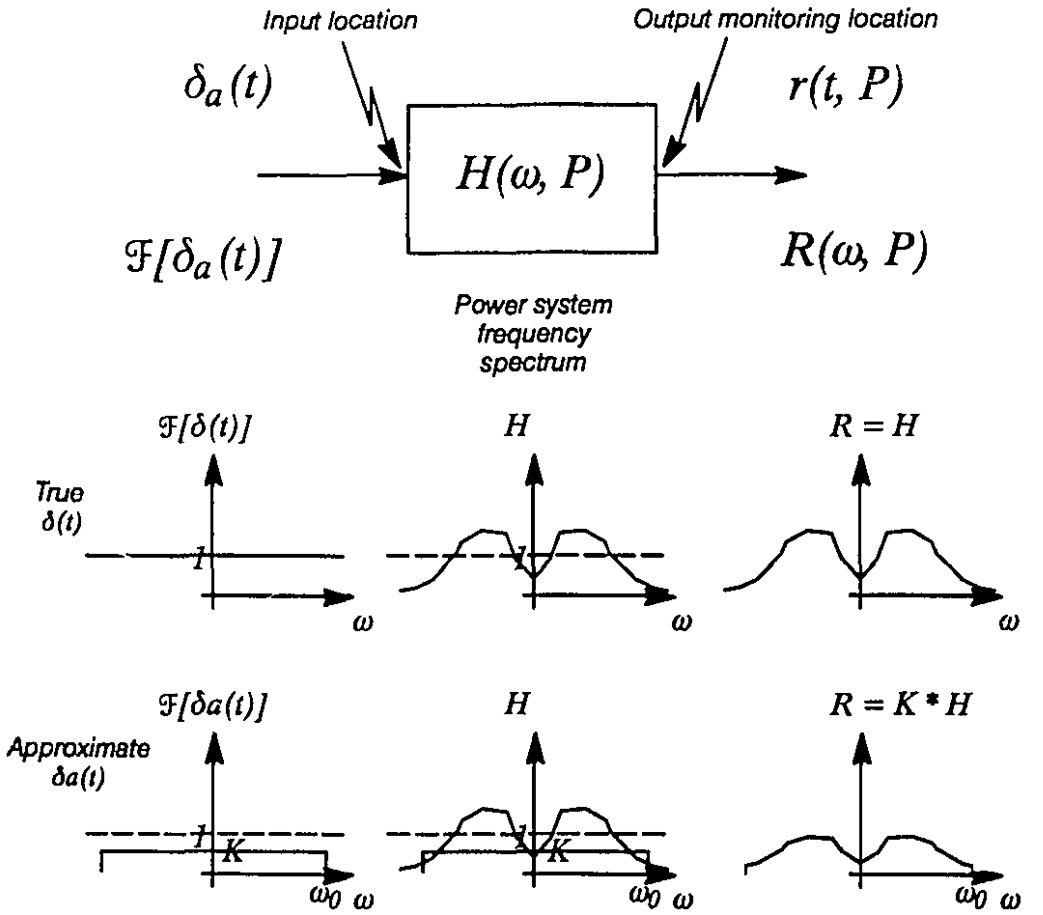

If the approximate impulse has:

1) *a known area K < 1*
2) *a relatively flat bandwidth with*
3) *a relatively high cutoff frequency ω_0 with respect to the system*

a good approximation of the impulse response is obtained.

Fig. 5–2. Estimating the network impulse response: The power system viewed as a filter.

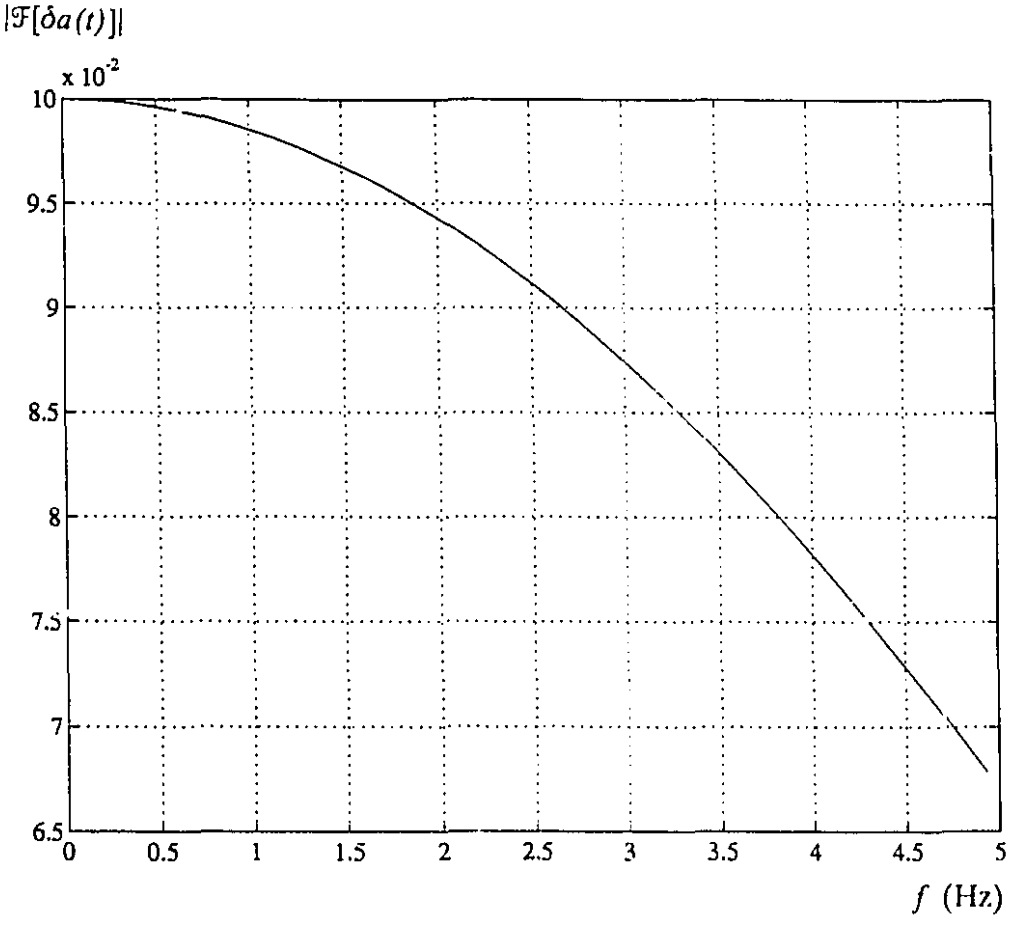

Fig. 5–3. Frequency spectrum of a 6–cycle, 3–phase fault.

pulse applied at the input.

It is well–known that the Fourier transform of a "true" unit impulse $\delta(t)$ equals 1, which also corresponds to its area [Frederick & Carlson 1971]:

$$\mathcal{F}[\delta(t)] = 1 \qquad (5-3)$$

Substituting (5–3) into (5–2), the transfer function $H(\omega, P)$ in this case is identically equal to the resulting output response frequency spectrum $R(\omega, P)$:

$$H(\omega, P) = R(\omega, P) \qquad (5-4)$$

An impulse may be approximated provided that the bandwidth of its frequency spectrum ω_0 is demonstrably larger than that of the filter to which it is applied, and that it is essentially flat over the same range. In the present case, the power system can be viewed as the filter: The requirement is therefore that the bandwidth of the approximate impulse be larger than that of the power system frequency spectrum, as illustrated in Fig. 5–2.

In our case, the proposed impulse is a 6–cycle, 3–phase "fault pulse" whose amplitude corresponds to a 1.0 p.u. voltage drop at t = 0. From the $v(t, P)$ frame of reference, this change corresponds to a momentary *absence* of voltage from 1.0 p.u. However, from the perspective of the rest state, the system cannot be *physically hit* by a pulse of greater amplitude.

Figure 5–3 shows the Fourier transform of such a pulse, obtained using the MATLAB software [Little & Shure 1992], which corresponds, as one would expect, to a sinc function: This would be even more apparent if the contribution of higher frequency elements were included. Fig. 5–4 shows the Fourier transform of a typical Hydro–Québec system transient response as defined in equation (5–1): These particular frequency spectra were determined for the Duvernay waveforms of Fig. 5–1 corresponding to the base case and stability limit (i.e. base+350 MW).

Comparing these frequency spectra, we observe that frequency component contributions can be considered negligible beyond 2 Hz on the Hydro–Québec

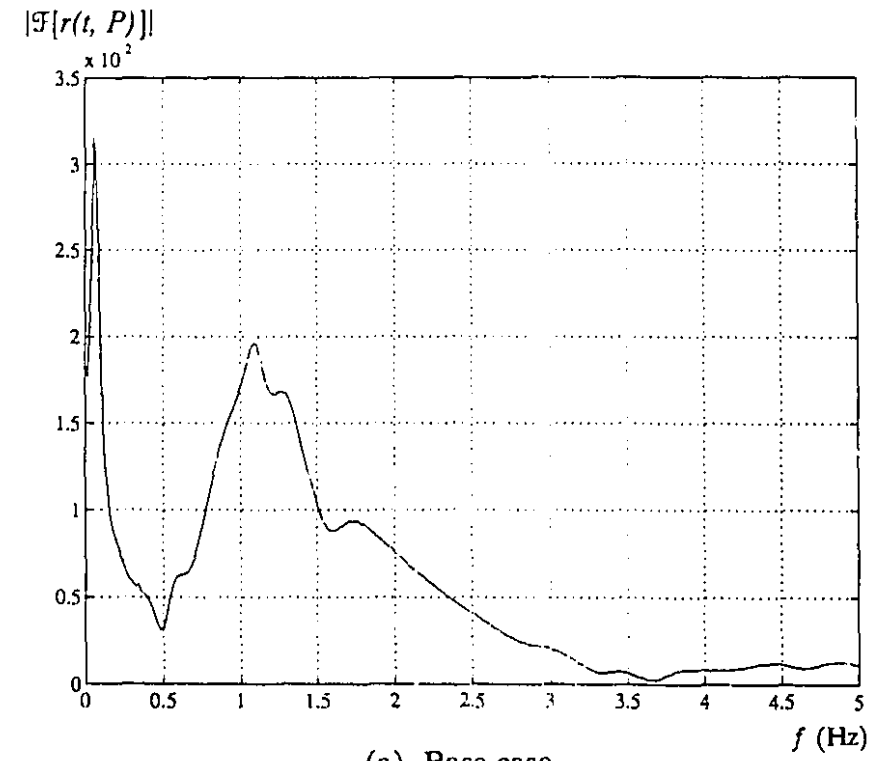

(a) Base case

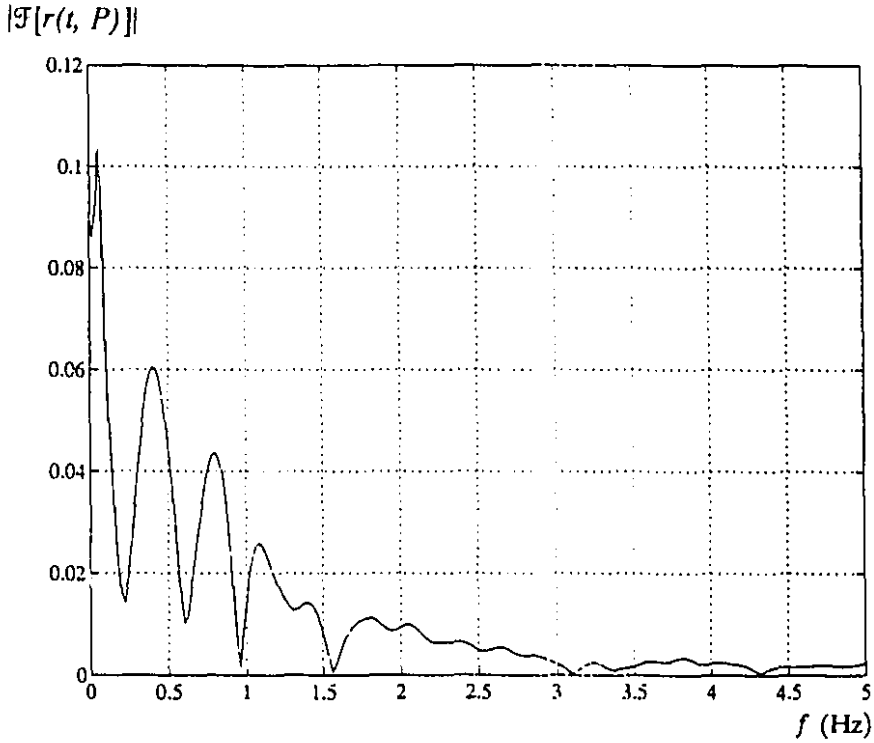

(b) Base case + 350 MW

Fig. 5–4. Frequency spectra of the transient response $r(t, P)$ of the Hydro–Québec network for two different values of power transfer. Monitored at Duvernay following a 6–cycle, 3–phase fault applied at Abitibi (3000 cycle total simulation time).

network (Fig. 5−4), whereas the pulse spectrum shows only a 6% drop over this range (Fig. 5−3). Additionally, one calculates the area under the 6−cycle, 3−phase fault to be (1 p.u. voltage x 0.1 seconds =) 0.1 p.u.−s., and this is confirmed by the value found at 0 Hz on Fig. 5−3. One concludes that a 6−cycle, 3−phase fault is therefore a good approximation of an impulse, within some factor K:

$$\mathcal{F}[\delta_a(t)] \simeq K \qquad (5-5)$$

where K = 0.1 over the 2 Hz frequency range for a 6−cycle pulse.

Substituting (5−5) into (5−2), the transfer function is now approximately given by:

$$H(\omega, P) \simeq R(\omega, P) / K \qquad (5-6)$$

Performing the inverse transform to return to the time domain, one obtains:

$$h(t, P) \simeq r(t, P) / K \qquad (5-7)$$

Expressing this in another way, an approximate network impulse response $h_a(t, P)$ can be defined as follows:

$$h_a(t, P) = r(t, P) / K \qquad (5-8)$$

To summarize, provided that the approximate impulse respects the conditions outlined in Fig. 5−2, one need not perform Fourier transforms of the input and output followed by an inverse transform to obtain the impulse response. In such instances, the transient response $r(t, P)$ to a 6−cycle, 3−phase fault is a faithful reproduction of a power system's impulse response under different initial operating conditions, apart for some linear scaling factor K dependent on fault duration.

5.3 TRANSIENT RESPONSE SIGNAL ENERGY AND ITS CORRELATION TO CONTINGENCY SEVERITY

From Fig. 5−1, we observe that steady−state generation prior to the contingency has considerable impact on the amplitude of the response to the

above–defined impulse: The higher the generation, the greater the amplitude of the response (i.e. the *swing)* until it degenerates into voltage collapse, corresponding to system instability (see Chapter 2). Intuitively, one might speculate that some measure of the swing may yield information regarding the system's approach to instability. One such measure is the signal energy of the impulse response [Frederick & Carlson 1971], defined as follows:

$$S(P) = \int_{-\infty}^{\infty} h(t, P)^2 \, dt \qquad (5-9)$$

To apply this definition to simulations of finite duration, one need only assume that, to all intents and purposes:

$$h(t, P) = 0, \quad t < t_f \text{ and } t > t_s \qquad (5-10)$$

where t_f is the fault (i.e. impulse) duration and t_s is the total simulation time. In Fig. 5–1, t_f = 6 cycles and t_s = 600 cycles (i.e. 10 seconds). Substituting (5–8) into (5–9), impulse response signal energy is therefore approximated by:

$$S(P) \simeq \frac{1}{K^2} \int_{-\infty}^{\infty} r(t, P)^2 \, dt \qquad (5-11)$$

As the integral involves only the transient response, let us ignore the constant term and focus on the signal energy of the transient response:

$$E(P) = \int_{-\infty}^{\infty} r(t, P)^2 \, dt \qquad (5-12)$$

The approach is summarized in Fig. 5–5. The points plotted on Fig. 5–6 illustrate typical results. These show the signal energy of the network transient response as a function of the total James Bay power generation (i.e. the parameter *P*) for three different stability limit searches. The signal energy results shown for a fault at Abitibi were obtained from the time–domain voltage waveforms of Fig. 5–1. In all cases, the output is taken as the voltage at Duvernay and signal energy is obtained by direct numerical integration of the transient stability ouput waveforms using equations (5–1) and (5–12).

Because voltage is in per unit (p.u.), signal energy is given in units of *p.u.*^{2}*–seconds* and power in the La Grande complex is given relative to the base

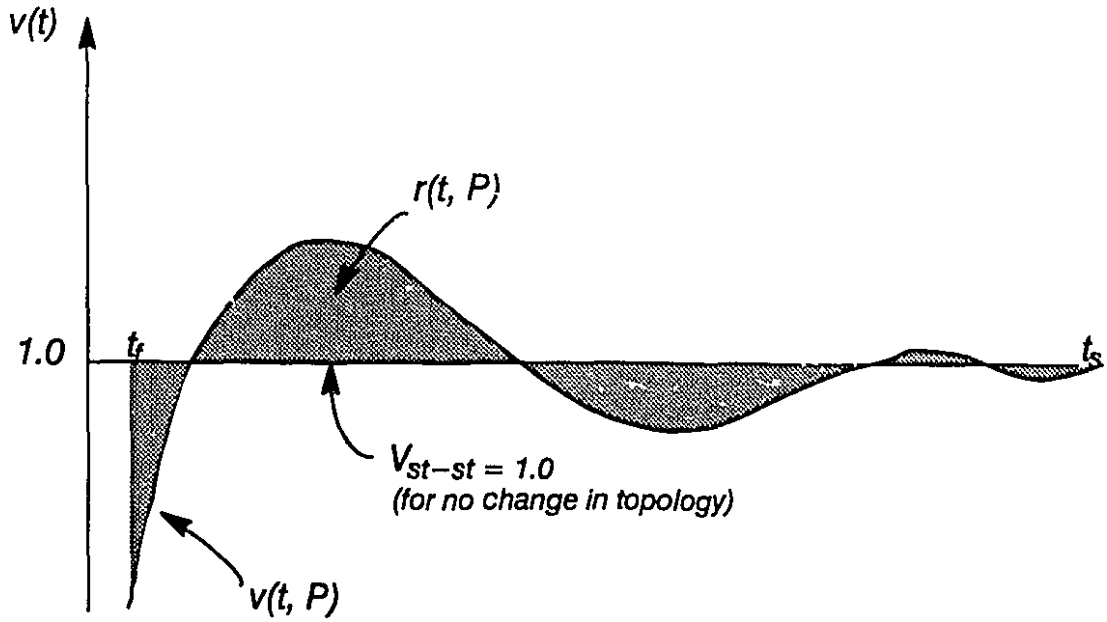

Define transient response as:

$$r(t, P) = v(t, P) - 1.0$$

Transient response signal energy is given by:

$$E(P) = \int_{-\infty}^{\infty} r(t, P)^2 \, dt \simeq \int_{t_f}^{t_s} r(t, P)^2 \, dt$$

Fig. 5–5. Transient response signal energy: Faults with no change in topology.

case of 8830 MW (i.e. 0 MW on the plot) from which all limit searches were initiated. If we consider the highest signal energy value on this plot for a fault at Abitibi, i.e. 52.3 x 10^{-4} $p.u.^2-s.$, we see that it occurs 350 MW above the base case (i.e. at 9180 MW) and that the stability limit corresponds to a simulation 25 MW above this value.

Considering how the Duvernay voltage waveforms are progressively distorted as the network is driven towards instability, regardless of contingency location, the smooth, monotonic rise of signal energy is a remarkable result. This is all the more so that a large number of active voltage–support devices are employed in the simulations (i.e. generators, synchronous condensers, SVCs, static excitation systems, etc.). For a given fault and location, the impulse or transient response's signal energy appears to quantify contingency severity as a function of increasing generation or power transfer. In addition, the asymptote is clearly identifiable as the stability limit for the particular contingency which, in turn, characterizes the severity of one contingency with respect to another.

A final point. If one imposes a typical acceptability criterion to the time–domain waveforms of Fig. 5–1 (e.g. voltage greater than 0.95 p.u. after 300 cycles) and observes where these are found on the signal energy plots of Fig. 5–6, acceptable cases are generally found in the region of the knee, where signal energy just begins to rise rapidly.

Let us now address the issue of explaining signal energy behaviour near the stability limit.

5.4 A LINEAR MODEL FOR SIGNAL ENERGY AS A NETWORK NEARS INSTABILITY

From a "black box" point of view, if a system's impulse response $h(t)$ is known, its transfer function $H(\omega)$ can be obtained by applying the Fourier transform [Frederick & Carlson 1971]:

$$H(\omega) = \int_{-\infty}^{\infty} h(t)\, e^{-j\omega t}\, dt \qquad (5-13)$$

For systems whose transfer function is not explicitly known, an approximate impulse response can be obtained by exciting the input using some appropriate

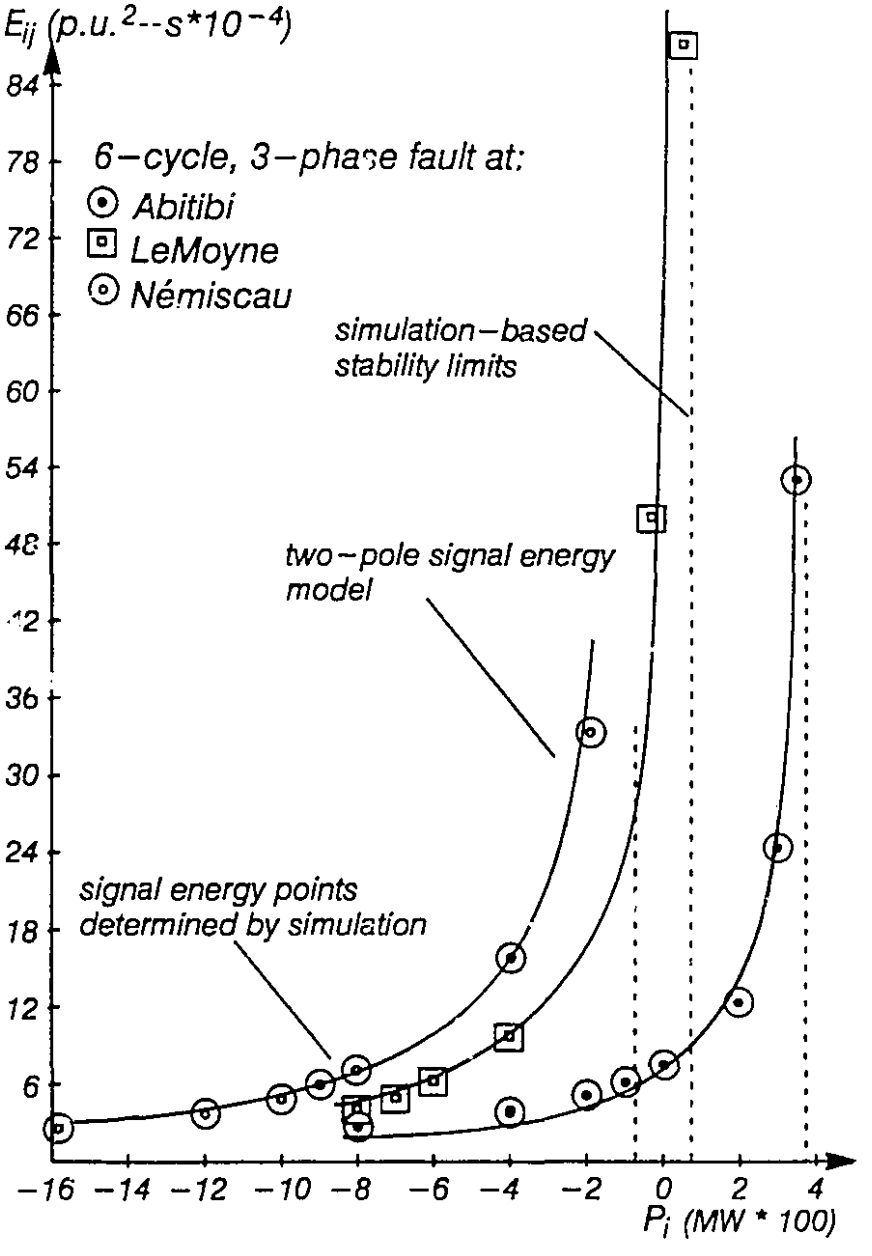

Fig. 5−6. Transient response signal energy behaviour at different fault locations as a function of the combined generation of LG 2, 3 and 4 (P_i). P_i is given relative to the base case *(0 MW)* corresponding to 8830 MW. For 6−cycle 3−phase fault with no loss of line.

impulse; then, performing a numerical integration of (5−10), one obtains an estimate of $H(\omega)$. This approach is equally valid whether $h(t)$ results from field measurements or from simulation. From this starting point, a great deal can be learned about a system using the large body of techniques available in the area of Fourier and Laplace transforms [Papoulis 1962], not to mention the fundamental work on frequency−domain and stability analysis contributed by Bode, Nyquist, Routh and others [Saucedo & Schiring 1968].

Of course, one may argue that such techniques, based on linear analysis, would only be marginally useful in such areas as modern power systems where we find many nonlinear elements such as generators, synchronous condensers, SVCs, static excitation systems, etc. However, as illustrated by Saitoh and Toyoda [Saitoh & Toyoda 1991], such techniques can *guide* us in our understanding of underlying physical phenomena in terms that are familiar and well documented (i.e. the frequency domain), and models built using such techniques can be extremely useful provided that we define their range of validity. On this basis, let us attempt to derive a formula which explains transient response signal energy behaviour using linear systems theory.

5.4.1 Model Derivation

From the perspective of transfer functions, stability theory tells us that a system is stable provided that all of its poles are on the left hand side of the complex plane [Frederick & Carlson 1971]. If one modifies some parameter causing a pole to move towards the imaginary axis, this reduces the damping (i.e. the real component) associated with the pole. In the time domain, reduced damping translates into greater amplitude when observing the transfer function's time−domain impulse response. As we have seen, this is precisely the type of behaviour observed in transient stability transfer limit searches of the kind shown in Fig. 5−1 where power generation is the parameter. Of course, for a system to be unstable, only a single right hand plane (RHP) pole is required.

Let us now consider the following. A real network may have thousands of poles and zeroes. As generation is changed, the impact on the network transfer function is to move a number of poles towards the $j\omega$ axis. When this happens, the first poles closing−in on the $j\omega$ axis tend to dominate all other contributions to the impulse response, and it becomes reasonable to express the transfer func-

tion in terms of only dominant poles. Since real systems showing damped oscillatory behaviour require the existence of at least two complex conjugate poles [Saucedo & Schiring 1968], p_{ij} and p_{ij}^*, the network transfer function $H_{ij}(\omega)$ for an impulse applied at location i and monitored at some predetermined location j can thus be represented as follows:

$$H_{ij}(\omega) = \frac{B_{ij}}{(j\omega - p_{ij})\,(j\omega - p_{ij}^*)} \qquad (5-14)$$

where

$$p_{ij} = \sigma_{ij} + j\omega_{ij}, \quad \sigma_{ij} < 0, \quad \omega_{ij} > 0 \qquad (5-15)$$

and B_{ij} is a constant while p_{ij} and p_{ij}^* represent the first pair of conjugate poles approaching the imaginary axis (i.e. the dominant poles), σ_{ij} and ω_{ij} respectively representing dominant pole damping and frequency.

Equation (5−6) tells us that the network transfer function $H_{ij}(\omega)$ differs from the transient response frequency spectrum only by a linear scaling factor K dependent on pulse duration t_f. The transient response frequency spectrum $R_{ij}(\omega)$ is therefore approximated by:

$$R_{ij}(\omega) = \frac{K_{ij}}{(j\omega - p_{ij})\,(j\omega - p_{ij}^*)} \qquad (5-16)$$

where p_{ij}, σ_{ij} and ω_{ij} are as above, and

$$K_{ij} = K B_{ij} \qquad (5-17)$$

We now wish to calculate the signal energy E_{ij} of the transient response $r_{ij}(t)$ in terms of $R_{ij}(\omega)$. One need only recall that Parseval's theorem [Taub & Schilling 1971] (sometimes known as Rayleigh's energy theorem [Frederick & Carlson 1971]) for real, energy−type signals bridges the gap between time and frequency domains by stating that:

$$E_{ij} = \int_{-\infty}^{\infty} r_{ij}(t)^2\, dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |R_{ij}(\omega)|^2\, d\omega \qquad (5-18)$$

where

$$|R_{ij}(\omega)|^2 = R_{ij}(\omega)\, R_{ij}^*(\omega) \qquad (5-19)$$

Combining equations (5−16) and (5−19) and rewriting:

$$R_{ij}(\omega)\, R_{ij}^*(\omega) = \frac{|K_{ij}|^2}{(\omega + jp_{ij})(\omega - jp_{ij})(\omega + jp_{ij}^*)(\omega - jp_{ij}^*)} \qquad (5-20)$$

To integrate the improper integral (5−18) using (5−20) as the integrand, one uses the calculus of residues [Churchill, Brown & Verhey 1974] and integrates along a contour encompassing the positive half of the complex plane (including the real axis). In this particular case, as $\sigma_i < 0$, the positive half complex plane for integrating purposes happens to be along the $-j\sigma$ axis, thus encircling the $-jp_{ij}^*$ and $-jp_{ij}$ poles, yielding the following theoretical result:

$$E_{ij}(\sigma_{ij}) = \frac{-|K_{ij}|^2}{(4\,\sigma_{ij})\,(\sigma_{ij}^2 + \omega_{ij}^2)} \qquad (5-21)$$

Note that E_{ij} is now explicitly expressed as a function of damping, σ_{ij}. Note also that E_{ij} is always positive as $\sigma_{ij} < 0$.

5.4.2 TRANSIENT RESPONSE SIGNAL ENERGY AS A FUNCTION OF POWER GENERATION

We have seen that, when searching for a transient stability transfer limit, generation P_i is the parameter used to drive the system to instability. Because the particular generation used to perform a limit search depends on fault location alone, and not monitoring location, only the subscript i need qualify this parameter.

In Fig. 5−1, we see that an increase in P_i clearly causes an increase in post−fault voltage swing, reflecting a decrease in system damping σ_{ij}. This phenomenon appears to accelerate as P_i approaches a finite value corresponding to the stability limit at $\sigma_{ij} = 0$, where the dominant poles intersect the imaginary axis, as sketched in Fig. 5−7. Based on this empirical evidence, it is reasonable to approximate damping σ_{ij} as a polynomial in P_i. For simplicity, near the stability limit, we consider the following first−order approximation:

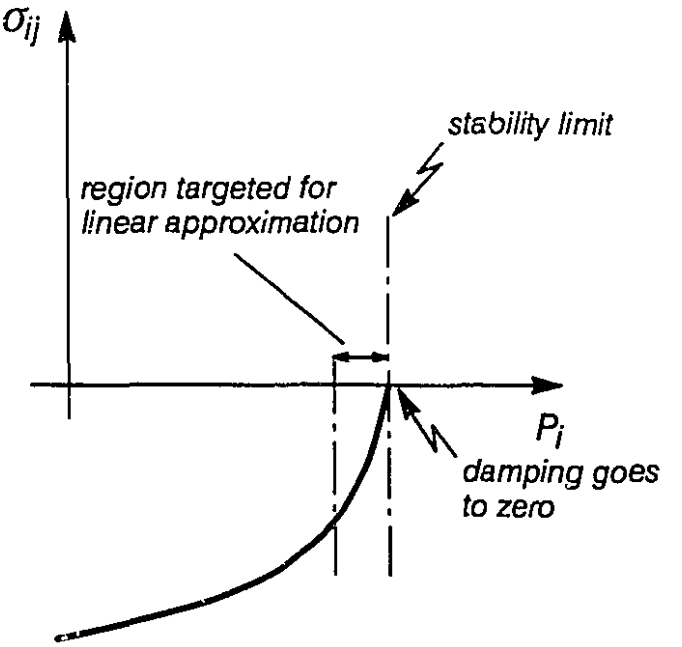

Fig. 5–7. Sketch of observed dominant pole damping (σ_i) behaviour as a function of generation P_i.

$$\sigma_{ij} = k_{ij} (P_i - C_{2ij}) \qquad (5-22)$$

where C_{2ij} is the constant maximum value of P_i at which the dominant poles cross over the $j\omega$ axis.

From Fig. 5–1, one notices that an increase in P_i also contributes to a decrease in ω_{ij}. This can be seen more clearly by comparing the frequency spectrum of the base case transient response (obtained using MATLAB) to that of the base+350 MW case (i.e. at the limit) shown in Fig. 5–8 for the voltage waveforms at the point of fault (i.e. at Abitibi, see Fig. 5–1 (b)): Increasing power by 350 MW causes dominant pole frequency to decrease by 20% whereas dominant pole damping decreases by 66% (i.e. amplitude increases by a factor of nearly 3). If these same cases are viewed from Duvernay (see Fig. 5–4), dominant pole frequency is seen to decrease by a factor of 2.75 whereas the amplitude again increases by a factor of 3. Though changes in P_i clearly impact ω_{ij}, the damping component σ_{ij} systematically appears to be more sensitive to such

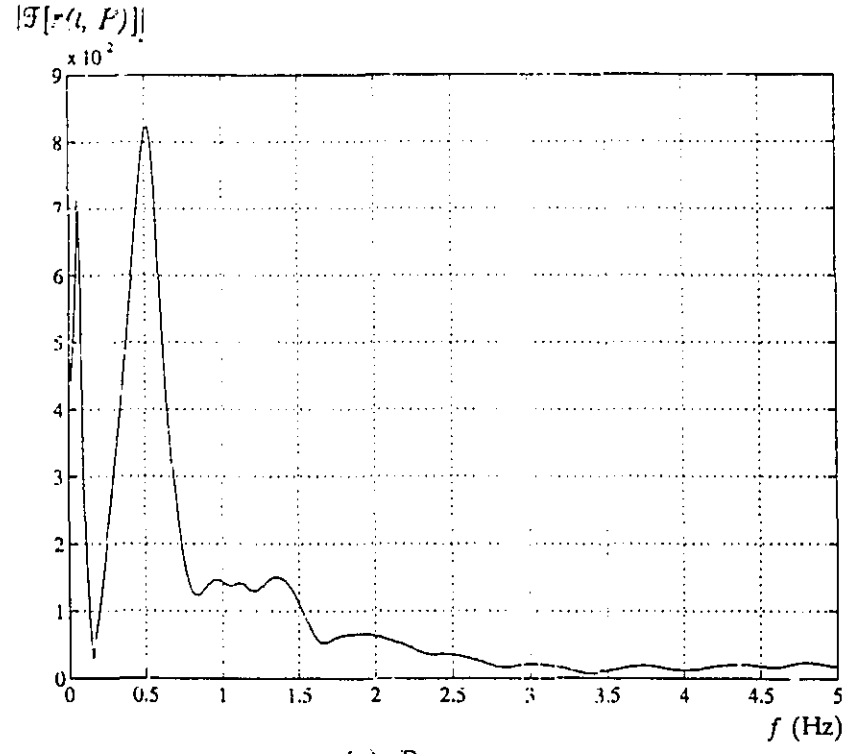

(a) Base case

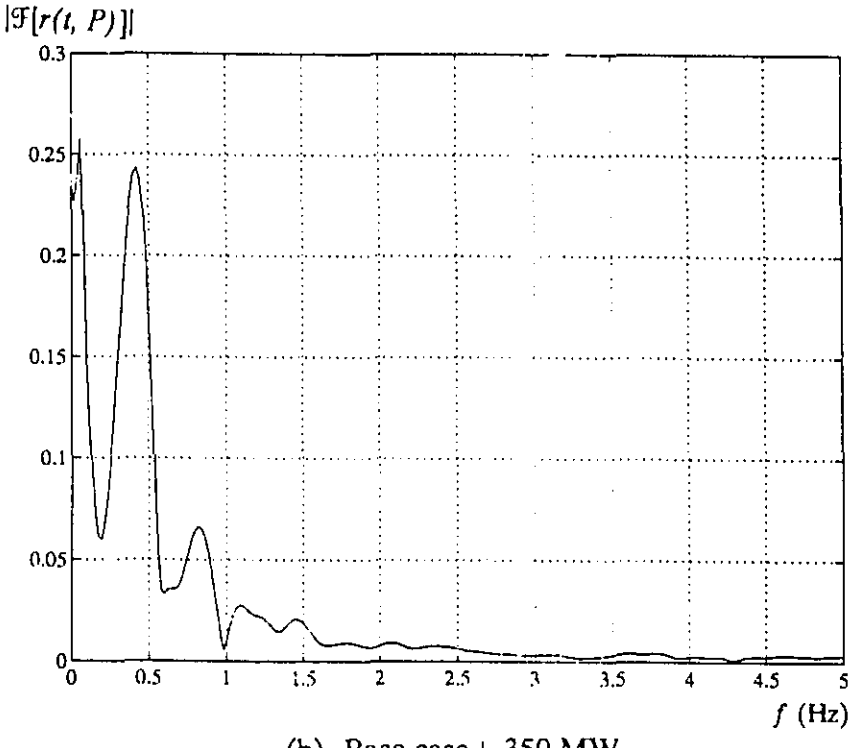

(b) Base case+ 350 MW

Fig. 5–8. Frequency spectra of the transient response $r(t, P)$ of the Hydro–Québec network monitored at Abitibi to a 6–cycle, 3–phase fault applied at Abitibi (3000 cycle total simulation time).

changes. Hence, changes in ω_{ij} are ignored in the present treatment.

Substituting (5−22) for σ_{ij} in equation (5−21), one obtains the following expression for the signal energy function near the crossover point:

$$E_{ij}(P_i) = \frac{-C_{1ij}}{(P_i - C_{2ij})\left[(P_i - C_{2ij})^2 + C_{3ij}\right]} \qquad (5-23)$$

where E_{ij} is now explicitly a function of the parameter P_i and

$$C_{1ij} = \frac{\pi |K_{ij}|^2}{2k_{ij}^3} \qquad (5-24)$$

$$C_{3ij} = \frac{\omega_{ij}^2}{k_{ij}^2} \qquad (5-25)$$

As the system approaches instability, $\sigma_{ij} \to 0$, implying that $(P_i - C_{2ij}) \to 0$. Under these circumstances, equation (5−23) can be simplified as follows:

$$\begin{aligned} E_{ij}(P_i) &\simeq \frac{-C_{1ij}}{C_{3ij}} \frac{1}{(P_i - C_{2ij})} \\ &= -C_{0ij} \frac{1}{(P_i - C_{2ij})} \end{aligned} \qquad (5-26)$$

Two important observations can be made. First, equation (5−26) tells us that signal energy E_{ij} of the transient response varies inversely with power near the stability limit. Second, the equation explicitly gives this limit, which happens to be the coefficient C_{2ij}. This coefficient defines the asymptote of equation (5−26) where $E_{ij}(P_i)$ theorically rises to infinity. Let us now compare equations (5−23) and (5−26) to simulation results obtained using ELISA.

5.5 Comparison of Model and ELISA Simulation Results

5.5.1 Stability Limit Estimation

The points plotted on Fig. 5−6 represent signal energy values obtained by simulation whereas the curves show the behaviour of equation (5−26) for coefficients obtained from two sample simulations.

5.5.1.1 SIGNAL ENERGY CURVES

For example, the curve for a fault at Némiscau was obtained using the two points indicated in Table 5−1. These points, separated by 400 MW, were chosen in the region of the knee (i.e. at a distance of 3.7% and 8.3%, or 325 and 725 MW respectively from the limit) in order to find two coefficients (i.e. C_{2i} and C_{0i}) for the purpose of *extrapolating* towards the stability limit. The Abitibi and LeMoyne curves were also obtained using the two points identified in Table 5−1 (again 400 MW apart), one of which was extremely close to the stability limit in each case (i.e. at a distance of .03% and .08%, or 25 and 75 MW respectively from the limit), in order to observe the effect of *interpolating* between points. A 400 MW distance between points was used at all times in order to verify equation (5−26)'s limit−estimating capability with the relatively coarse initial steps used by ELISA in its present search strategy.

In all cases, a good fit is observed between the model curves and simulation results, particularly in the region beyond the knee, close to the stability limit. These results show conclusively that relatively distant poles have far less effect on signal energy behaviour than the first pair of poles (i.e the dominant poles) closing−in on the imaginary axis.

5.5.1.2 SIGNAL ENERGY LIMIT ESTIMATES

A practical aspect of equation (5−26) is to attempt to estimate stability limits from simulation results and see how these values differ from limits obtained by exhaustive simulation. This also permits an objective evaluation of the validity of equations (5−23) and (5−26) in meaningful engineering terms. To calculate the coefficients, equation (5−26) requires only two sample simulations and the computations are simple; equation (5−23) requires three sample simulations and the computations require solving a system of nonlinear equations.

Table 5−1 shows the error between the stability limits calculated using equation (5−26) and those obtained by simulation using ELISA for 6−cycle, 3−phase faults applied at each of seven stations along the James Bay corridor: System output is systematically taken to be the voltage at Duvernay. The simulation results give the stability limits within 25 MW. For instance, if we consider the fault at Albanel when calculating the C_{2i} from two points set 400 MW apart, the closest being within 350 MW or 4% away from the known simulation−based

TABLE 5–1

Comparison of Estimated and True[1] Stability Limits For 6–cycle 3–phase Fault With No Loss of Line

(i.e. using points close to the limit)

$$E_{ij}(P_i) = \frac{-C_{0ij}}{(P_{ij} - C_{2ij})}$$

Fault Location	*Points[2] (p.u.)*		*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*
LG2	.9469	.9941	8495	8480	0.18
Chibougamau	.9541	.9973	9274	9255	0.21
Abitibi	.9538	.9973	9238	9205	0.36
LeMoyne	.9467	.9916	8927	8905	0.25
Némiscau	.9172	.9629	8755	8755	0
Albanel	.9146	.9601	8772	8780	0.10
Radisson	.9386	.9854	8574	8555	0.22

TABLE 5–2

Comparison of Estimated and True[1] Stability Limits For 6–cycle 3–phase Fault With No Loss of Line

(i.e. using distant points)

$$E_{ij}(P_i) = \frac{-C_{1ij}}{(P_i - C_{2ij})[(P_i - C_{2ij})^2 + C_{3ij}]}$$

Fault Location	*Points[2] (p.u.)*			*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*
LG2	.8526	.8998	.9469	8409	8480	0.84
Chibougamau	.8676	.9109	.9541	9311	9255	0.61
Abitibi	.8724	.9158	.9593	9310	9205	1.14
LeMoyne	.9055	.9467	.9916	8942	8905	0.42
Némiscau	.8715	.9172	.9629	8825	8755	0.80
Némiscau	.8258	.8715	.9172	8545	8755	2.40
Albanel	.8235	.8690	.9146	8554	8780	2.57
Radisson	.8451	.8919	.9386	8420	8555	1.58

1 *Limit obtained by simulation: within 25 MW*
2 *Points used to estimate limit, in p.u. of True Limit (all points separated by 400 MW)*
3 *% error with respect to stability limit obtained by simulation*

limit (and the farthest, 8.5%), one estimates a value within 8 MW or 0.1% of the stability limit.

Table 5−2 similarly compares stability limits calculated using equation (5−23) and those obtained by simulation, but for sample points farther from the limits. Using this equation, three points are required as three coefficients must be computed. The three points were once again chosen to be 400 MW apart, for the same reason as above. For example, considering again a fault at Albanel, the closest and the most distant points are respectively 750 and 1550 MW (i.e. 8.5% and 17.6%) away from the limit and equation (5−23) estimates a limit which is 226 MW more conservative than the simulated limit, within an accuracy of 2.6%. In general, Table 5−2 shows that the estimated limit improves as the set of sample points is taken closer to the limit: this is clearly seen from the two cases where a Némiscau limit is estimated from two different sets of points, one set being closer than the other by 400 MW.

5.5.2 Region of Validity of the Proposed Signal Energy Model for Faults with No Loss of Line

Figure 5−9 is identical to Fig. 5−6 except that the results are plotted on a per−unit basis with respect to the stability limit for each individual fault location. This figure illustrates the similarities in signal energy behaviour for faults at different locations, notwithstanding very different and, at times distorted, time−domain responses. Though, for clarity, only three signal energy plots are shown here, this holds for all other fault locations appearing in Tables 5−1 and 5−2.

In this figure, we see that signal energy appears to "explode" within approximately 5% of the stability limit. In this region, $(P_i - C_{2i})$ is less than 0.05 p.u., closely matching the assumption made in deriving equations (5−23) and (5−26) that it should be very small as the network approaches instability. This not only explains why, in Table 5−1, equation (5−26) gives high limit−estimating accuracy (less than 0.4% error) when the closest point is within 4% of the stability limit (though the furthest is within 8.5%) but also why, in Table 5−2, equation (5−23) gives an accuracy within only 2.6% when the closest point is within 8.5% (and the furthest, within 17.6%).

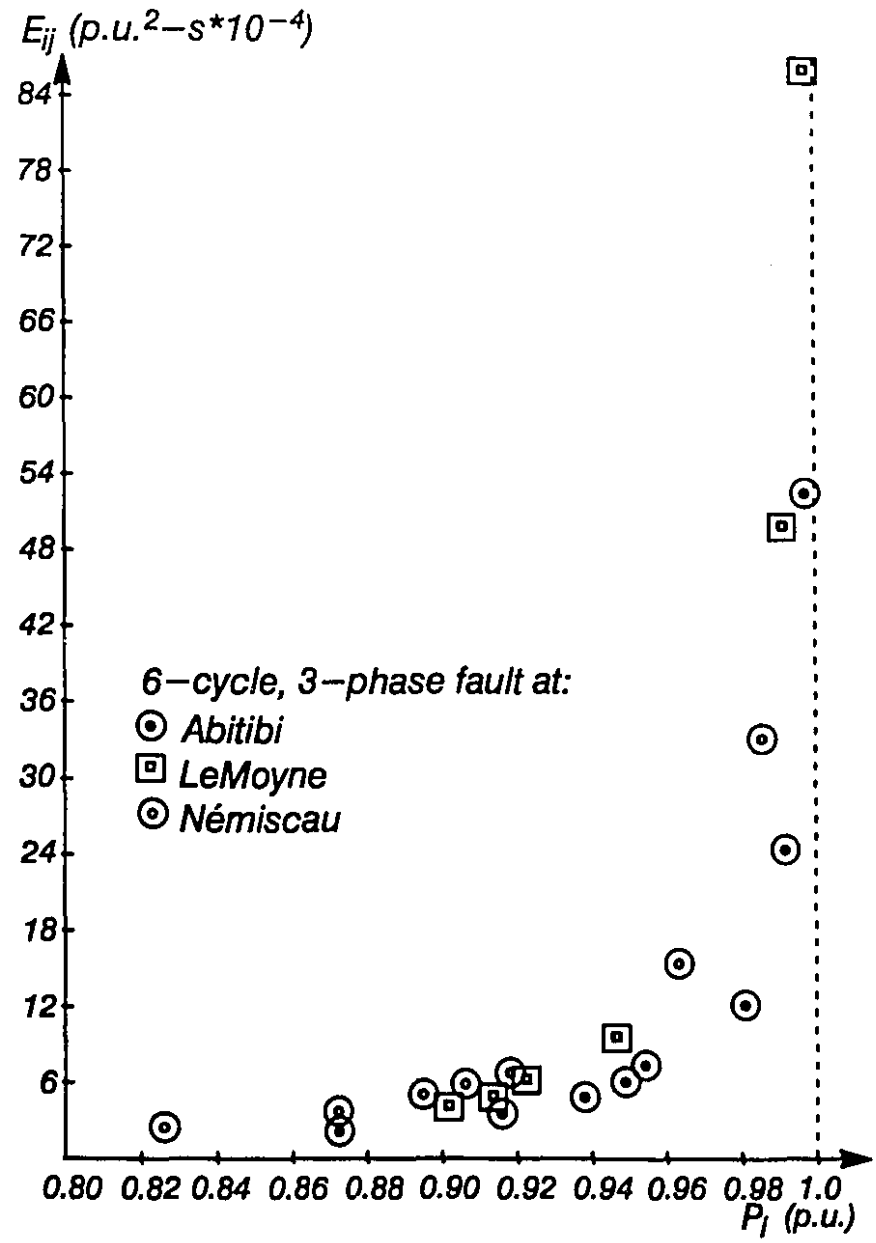

Fig. 5–9. Signal energy results of Fig. 5–6 plotted in per–unit of the simulated stability limit of each contingency (6–cycle 3–phase fault with no loss of line).

5.5.3 THE CONCEPT OF STABILITY MARGIN

Figure 5–9 also helps us understand the concept of stability margin, used almost universally by operations planners, and its positive effect on system stability. In particular, for faults with no subsequent change in topology, signal energy rises slowly up to about 95% of the stability limit. At this point, small increases in generation translate into large increases in signal energy. One clearly sees that the power system is far more robust at 95% of its stability limit than, for example, at 98%. This correlates with the observation, made earlier, that operating points exhibiting acceptable time–domain voltage behaviour appear to be on the knee, not on the fast–rising part of the signal energy plot.

5.5.4 THE EFFECT OF NON–LINEAR ACTIVE VOLTAGE SUPPORT

We have already mentioned that the Hydro–Québec network model used to perform ELISA simulations contains a large number of active, non–linear voltage–support components. Nonetheless, the proposed signal–energy model, based on linear analysis, has demonstrated considerable success in estimating stability limits, particularly when using at least one point beyond 95% of the limit.

We can offer some explanation for this phenomenon. At the stability limit, the active elements in the network – generators, SVCs, synchonous compensators, static excitations, etc. – have all reached their limit. This defines the stability limit: Nothing further can be done without the action of more equipment. The system at this point is static: It is fixed in time and space. It is not surprising, therefore, that the system is amenable to linear analysis and that the estimated limits thus obtained are accurate within tenths of a percent. By the same reasoning, limits obtained at 80–90% of the stability limit are only accurate within a few percent because the system at this stage is not completely static: The active elements are still at work to some extent, though far less than in the 90–96% region because the system is inherently more robust at great distance from the limit. Consequently, active, nonlinear voltage–support elements contribute to *translating the stability limit to a higher value by slowing the rate of rise of signal energy in the region of the knee.*

A corollary to this is that one should be capable of characterizing voltage support elements by their effect on signal energy curvature in the region of the

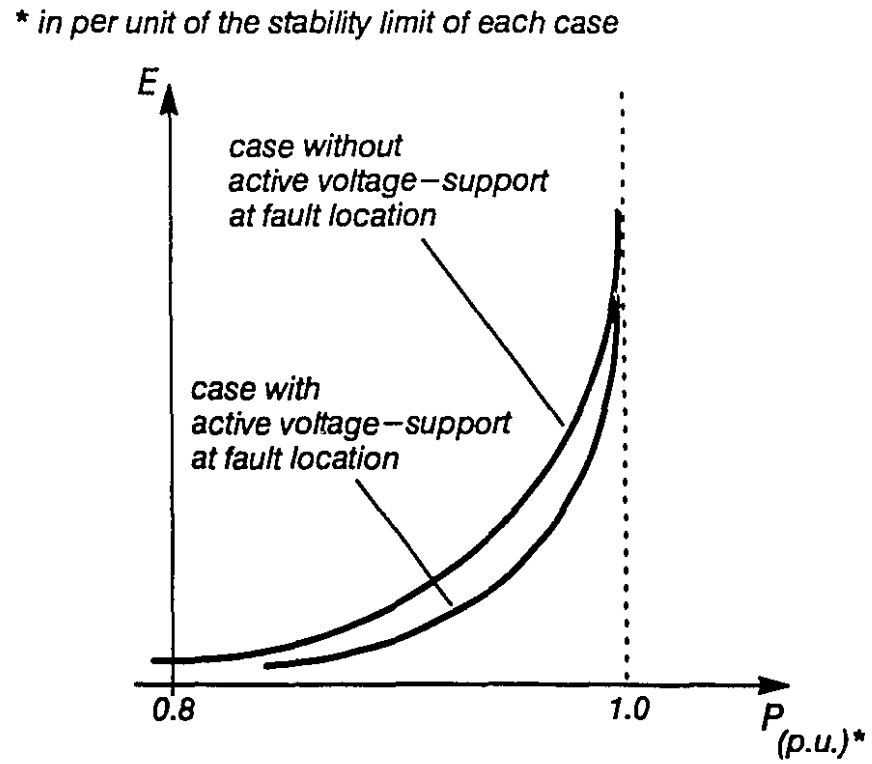

Fig. 5−10. Impact of voltage−support device on signal energy behaviour.

knee. If such elements are present at or near a particular fault location, the rise of the signal energy curve as a function of power in per unit of the stability limit should initially be slower than if there were none, and suddenly become sharper beyond a certain point (i.e. the point at which they are no longer capable of actively controling the rate of rise in signal energy), as sketched in Fig. 5−10. In fact, Fig. 5−9 exhibits this behaviour, the results for a fault at Némiscau (no SVCs, no SCs) showing a more gradual rise than those for a fault at Abitibi (SCs with static excitation) or LeMoyne (LG4 generators with static excitation).

5.6 NORMAL CONTINGENCIES

We have so far only considered an elementary contingency consisting of a 6−cycle, 3−phase fault with no loss of line. Though of conceptual value, the results are of limited use unless it can be shown that the signal energy approach is valid for all normal contingencies [Hydro−Québec 1990]. Specifically, it must be shown that the approach is valid for the three classes of normal contingencies:

1) Fault with no loss of element;
2) Fault with loss of element;
3) No fault, spontaneous loss of element.

Signal energy stability limit estimation has so far considered only the first class. We now consider the other two. As we shall see, the approach can in fact be extended to cover the entire class of normal contingencies.

5.6.1 GENERALIZING THE SIGNAL ENERGY APPROACH TO ALL NORMAL CONTINGENCIES

5.6.1.1 REST STATE AND TRANSIENT RESPONSE

From our previous linear SISO perspective (section 5.2), as power system topology inherently determines the power system response to external inputs (i.e. the transfer function), a *change* in topology necessarily modifies the observed dynamic behaviour, or transfer function. However, changes in topology, with or without faults, have a more disturbing impact: Post−contingency steady−state bus voltages differ from those of the pre−contingency starting state.

This last issue is critical: The previously−described signal energy analysis hinges on identifying an appropriate rest state with respect to which the transient response is defined. If system frequency were the monitored quantity, this issue would not be raised because post−contingency frequency, regardless of monitoring location, inevitably returns to its pre−contingency steady−state value of 1. p.u. (i.e. 60 Hz) due to speed governor action. In the case of voltage however, which is now the appropriate rest state: The 1.0 p.u. pre−contingency starting state value? Or the steady−state post−contingency value V_{st-st}?

Consider a network which loses an element, whether or not it is preceded by a fault. If this modified network settles to a steady state which is different from that of the starting state, we have no alternative but to accept the former as the rest state for the *changed* network: The new network has itself defined its rest state. The initial steady state must therefore be viewed as a disturbance−dependent, non−zero initial condition which contributes to the transient response with respect to the post−contingency rest state.

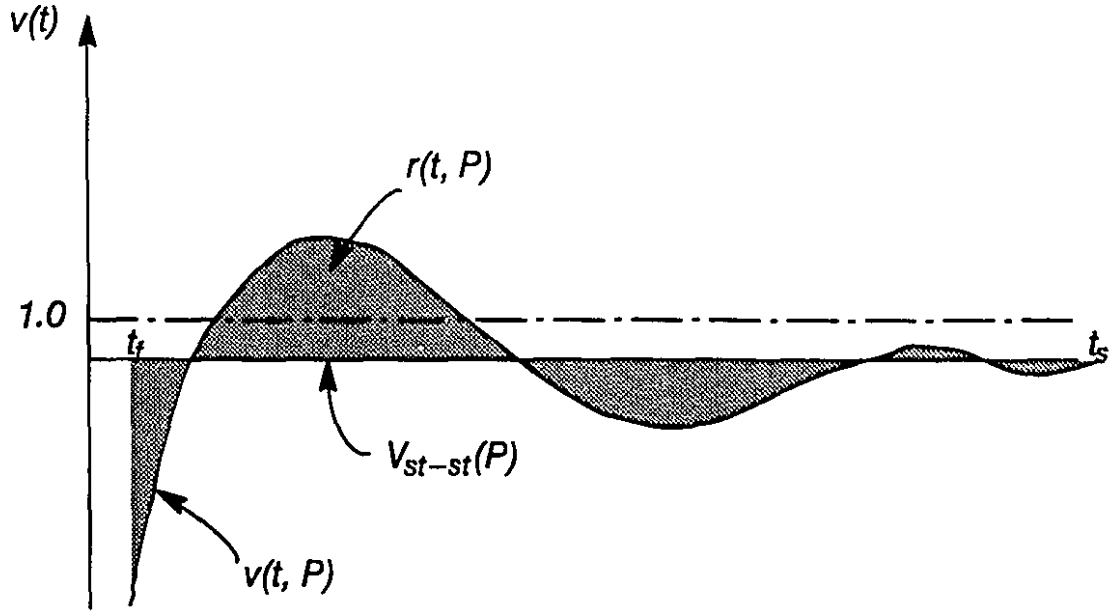

Define transient response as:

$$r(t, P) = v(t, P) - V_{st-st}(P)$$

Transient response signal energy is given by:

$$E(P) = \int_{-\infty}^{\infty} r(t, P)^2 \, dt \simeq \int_{t_f}^{t_s} r(t, P)^2 \, dt$$

Fig. 5–11. Transient response signal energy calculation for all normal contingencies. Post–contingency steady–state voltage is dependent on generation *P*.

Because post−contingency steady state voltage can be obtained *a priori* by means of appropriate power−flow analysis [IEEE 1991], without having to resort to transient stability software, the network transient response need only be defined as follows:

$$r(t, P) = v(t, P) - V_{st-st}(P) \qquad (5-27)$$

The post−contingency steady−state voltage is dependent on the operating point, as indicated in equation (5−27): This is particularly evident in the examples of Fig. 5−14. The signal energy calculation for normal contingencies is summarized in Fig. 5−11.

5.6.1.2 FREQUENCY DOMAIN MODEL

As we have mentioned before, a change in topology modifies the network transient response. In the frequency domain, such changes may be modelled as a disturbance function [Frederick & Carlson 1971] which acts on the system, thereby modifying network response as depicted in Fig. 5−12. In the general case where faults and disturbances are applied concurrently, the network transient response in the frequency domain is therefore given by:

$$R_{ij}(\omega, P) = \mathcal{F}[\delta_a(t)]\ H_{ij}(\omega, P) + D_{ij}(\omega, P) \qquad (5-28)$$

where i represents the location of the fault or disturbance, or the two combined, and j represents the monitoring location. Substituting (5−5) into (5−28), this simplifies to:

$$R_{ij}(\omega, P) \simeq K\ H_{ij}(\omega, P) + D_{ij}(\omega, P) \qquad (5-29)$$

Taking the inverse Fourier transform, we see that the transient response $r_{ij}(t, P)$ is simply a linear combination of the scaled−down impulse response $h_{ij}(t, P)$ and the disturbance time−domain component $d_{ij}(t, P)$:

$$r_{ij}(t, P) \simeq K\, h_{ij}(t, P) + d_{ij}(t, P) \qquad (5-30)$$

As mentioned before, it is useful to view the disturbance component $d_{ij}(t, P)$ as including the non−zero initial condition with respect to the post−contingency rest state in addition to the contribution due to a change in topology.

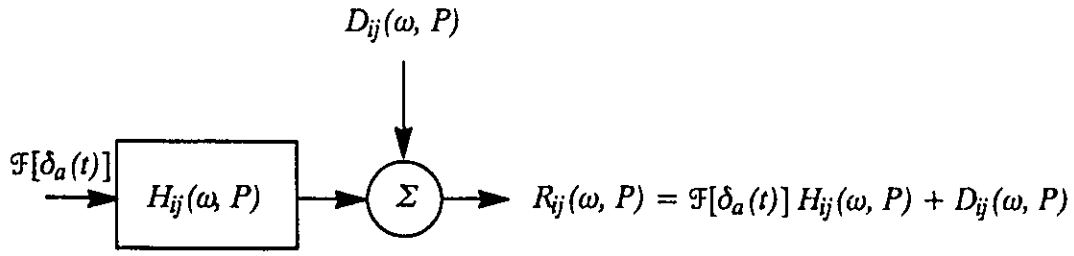

Fig. 5−12. Frequency domain transient response model for all normal contingencies.

Where no topology changes are included in the contingency, equation (5−29) reduces to:

$$R_{ij}(\omega, P) \simeq K H_{ij}(\omega, P) \tag{5−31}$$

for which the time response is given by:

$$r_{ij}(t, P) \simeq K h_{ij}(t, P) \tag{5−32}$$

This is essentially equation (5−7), the starting point for our previous analysis.

Where the contingency includes no fault, and a spontaneous change in topology occurs, equation (5−28) reduces to:

$$R_{ij}(\omega, P) = D_{ij}(\omega, P) \tag{5−33}$$

and the time response is given by:

$$r_{ij}(t, P) = d_{ij}(t, P) \tag{5−34}$$

The model therefore covers all normal contingencies. We now address the issue of deriving an appropriate signal energy model.

5.6.1.3 SIGNAL ENERGY MODEL FOR NORMAL CONTINGENCIES

For this new system to be unstable, a single RHP pole must be present in $H_{ij}'(\omega, P)$, corresponding to the system's response $R_{ij}(\omega, P)$ to a true unit impulse

(i.e. where $\delta_\alpha(t) = \delta(t)$ and $\mathcal{F}[\delta(t)] = 1$ in equation (5−28)):

$$H_{ij}'(\omega, P) = H_{ij}(\omega, P) + D_{ij}(\omega, P) \qquad (5-35)$$

Because a 6−cycle, 3−phase fault is a good approximation of an impulse, the transient response $R_{ij}(\omega, P)$ of a network to a normal contingency is expressed in terms of a linear combination of the system transfer function $H_{ij}(\omega, P)$, *essentially undistorted except for a scaling factor,* and an appropriate disturbance function $D_{ij}(\omega, P)$, as given by equation (5−29). The significance of this is that the poles of the new system transfer function $H_{ij}'(\omega, P)$ are *left intact* in the transient response frequency spectrum $R_{ij}(\omega, P)$: If the dominant poles of $H_{ij}'(\omega, P)$ approach the imaginary axis in the pole−zero plane due to increased power P, this behaviour can be observed, as before, by analysing that of the dominant $R_{ij}(\omega, P)$ poles. Consequently, the derivation of a signal energy function for all normal contingencies, employing equation (5−16) as a starting point, can be considered valid for faults including changes in topology, and equations (5−23) and (5−26) are therefore applicable to all normal contingencies. We now consider the use of these formulae for stability limit estimation.

5.6.2 STABILITY LIMIT ESTIMATION: FAULTS WITH LOSS OF LINE

Figure 5−13 shows signal energy as a function of power P for 6−cycle, 3−phase faults followed by the loss of a line section south of the faulted station for three fault locations. Once again, this figure clearly shows the signal energy's asymptotic approach to the stability limit for each contingency location, as in the case of Fig. 5−6 for 3−phase faults and no change in topology. The curves are calculated using the two points identified in Table 5−3 (and equation 5−23). Steady−state post−contingency voltage at the monitoring location is obtained by executing 3000−cycle simulations rather than performing a post−contingency power flow analysis. The limit−search results shown in Fig. 5−14 are typical in that the network has clearly settled to a steady state by this time.

As can be seen in Fig. 5−13, the 2−point formula (i.e. equation 5−26) results in a very good fit to the simulation−based signal energy points obtained for the faults at Chamouchouane and Némiscau: This is due to the fact that at least one of the points is quite close to the limit (see Table 5−3) and the farthest

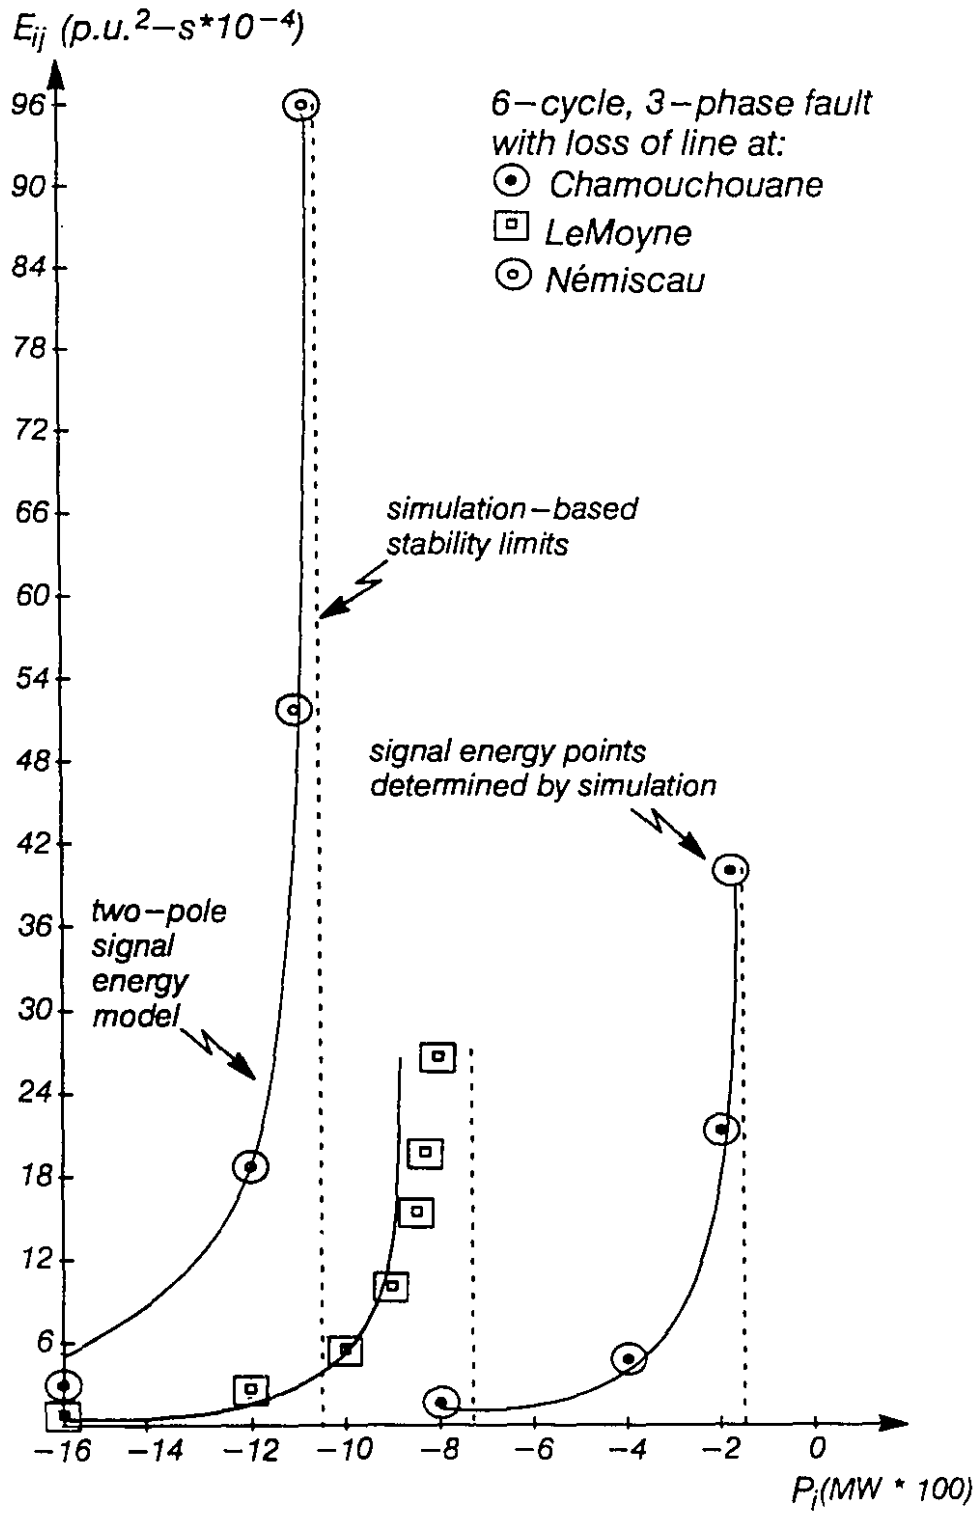

Fig. 5–13. Transient response signal energy behaviour at different fault locations as a function of the combined generation of LG 2, 3 and 4 (P_i). P_i is given relative to the base case (*0 MW*) corresponding to 8830 MW. For a 6–cycle, 3–phase fault with loss of line.

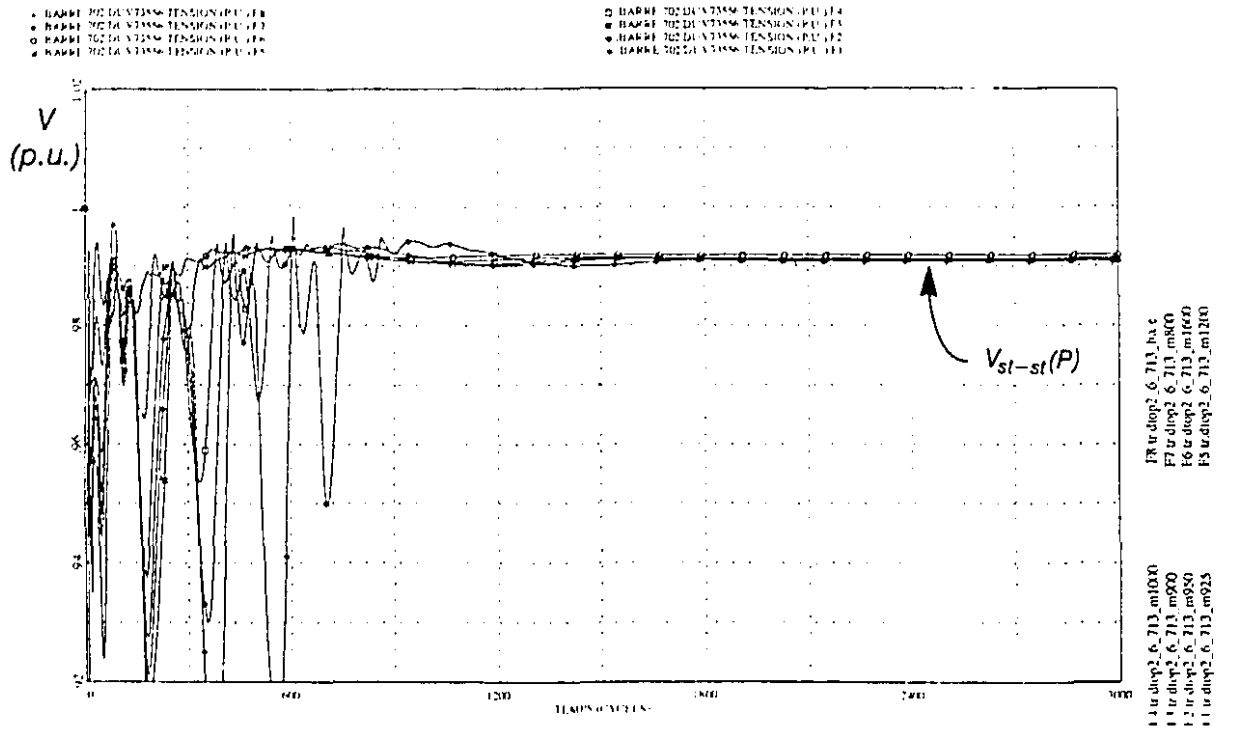

(a) Voltage monitored at Duvernay

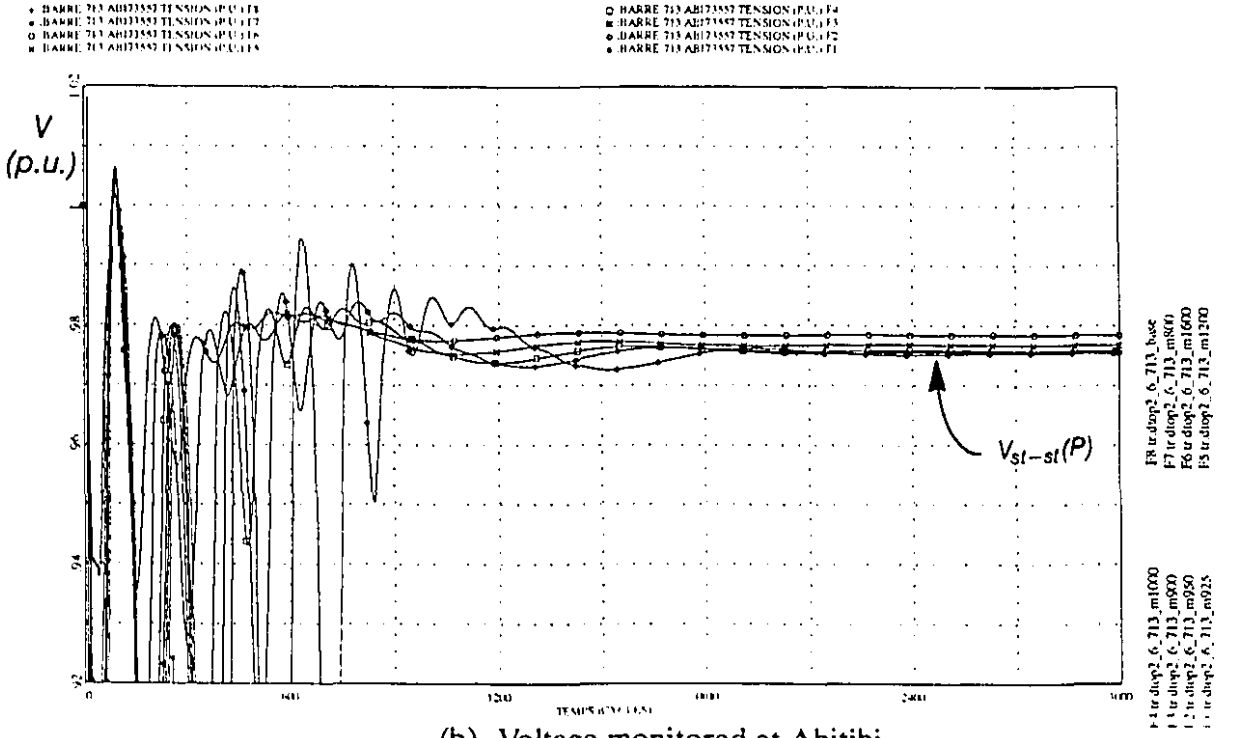

(b) Voltage monitored at Abitibi

Fig. 5−14. Transient stability transfer limit search for 6−cycle, 3−phase fault at Abitibi with loss of line. Note dependency of V_{st-st} on pre−contingency power transfer.

point is less than 10% away. The curve for the fault at Le Moyne, near the limit, coincides less with the points obtained by simulation and this is corroborated by the higher error obtained in the limit estimate (see Table 5–3): In this case, one observes that the farthest point is slightly greater than 10% from the limit.

Tables 5–3 and 5–4 once again illustrate the limit–estimating capability of equations (5–23) and (5–26). In Table 5–3 (i.e. 2–point formula), only two have an error greater than 1.5%. Of these two, the error for LaVérendrye is clearly due to the distance of the second point, this being 16% away from the limit. It is useful to note that, at Hydro–Québec, limits are found within 100 MW, which roughly corresponds to 1% on both Tables. The error in the Churchill Falls limit estimate will be treated in 5.6.4.

Table 5–4 shows results using the 3–point formula (equation 5–23). In this case, the same points are used, except that a third point is now added in between the previous two for each fault location. The error systematically drops in each case, except in that case where the error was already negligible (i.e. at Chamouchouane: 0.06%). Of the nine, a single case shows an error greater than 1% (i.e. once again at Churchill Falls: 2.36%). On the Hydro–Québec network, particularly the James Bay transmission system, it appears that less than 1% error can be expected in limit–estimating capability using the 3–point formula provided that at least one point is within 3.5% and the farthest within 16% of the limit. Had the 3–point formula been used, a close fit would have resulted for all of the contingencies illustrated in Fig. 5–13.

5.6.3 STABILITY LIMIT ESTIMATION: SPONTANEOUS LOSS OF LINE

Table 5–5 illustrates the limit–estimating capability of equations (5–23) and (5–26) for the spontaneous loss of a line, without fault. Two cases are presented using the 2–point formula (equation 5–26) whereas a single case is shown of the 3–point formula (equation 5–23).

One notices that the two–point formula is applied for the two combinations of the three points used in the three–point formula including the closest point (i.e. at 95.5% of the limit). When the two points are within 5% and 10% of the limit, equation 5–26 gives exceptionally good estimates as can be seen

TABLE 5–3

Comparison of Estimated and True[1] Stability Limits For 6–cycle 3–phase Faults With Loss of Line Using

$$E_{ij}(P_i) = \frac{-C_{0ij}}{(P_i - C_{2ij})}$$

Fault Location	*Points[2] (p.u.)*		*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*
LG2	.8602	.9792	8339	8405	0.79
Le Moyne	.8920	.9661	7990	8105	1.42
Némiscau	.9293	.9936	7760	7780	0.26
Albanel	.9293	.9807	7703	7780	0.99
Abitibi	.9117	.9874	7867	7930	0.79
Chibougamau	.9056	.9787	8102	8205	1.26
Chamouchouan	.9293	.9942	8685	8680	0.06
La Vérendrye	.8382	.9635	10010	9580	4.49
Churchill Falls	.8976	.9721	7845	8055	2.61

TABLE 5–4

Comparison of Estimated and True[1] Stability Limits For 6–cycle 3–phase Faults With Loss of Line Using

$$E_{ii}(P_i) = \frac{-C_{1ij}}{(P_i - C_{2ij})[(P_i - C_{2ij})^2 + C_{3ij}]}$$

Fault Location	*Points[2] (p.u.)*			*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*
LG2	.8602	.9554	.9792	8393	8405	0.14
Le Moyne	.8920	.9414	.9661	8037	8105	0.84
Némiscau	.9293	.9807	.9936	7784	7780	0.05
Albanel	.9293	.9550	.9807	7731	7780	0.64
Abitibi	.9117	.9622	.9874	7922	7930	0.10
Chibougamau	.9056	.9543	.9787	8135	8205	0.85
Chamouchouan	.9293	.9712	.9942	8693	8680	0.15
La Vérendrye	.8382	.9217	.9635	9662	9580	0.86
Churchill Falls	.8976	.9472	.9721	7865	8055	2.36

[1] *Limit obtained by simulation: within 25 MW*
[2] *Points used to estimate limit, in p.u. of True Limit*
[3] *% error with respect to stability limit obtained by simulation*

in the example shown. However, if the second point is too far away (i.e. at 19% from the limit), limit−estimation capability is severely degraded, as one would expect with the simplified formula. Nonetheless, when the three−point formulat is used, the same three points estimate a limit whose error is well within 1%.

5.6.4 THE EFFECT OF MONITORING LOCATION

For purposes of comparison, the signal energies were systematically computed for the same monitoring location, this being the 735 kV station at Duvernay. Though the issue of monitoring location, including that of multiple locations, remains to be studied in depth, results have been monitored at different locations and a number of observations can be made:

a) Signal energy for stability limit estimation can be monitored at any point in the EHV network where voltage can be tapped.
b) Signal energy values monitored at different locations for the same contingency will be different, but they will rise asymptotically to the same stability limit.
c) Signal energies calculated from results obtained at monitoring locations close to that of the contingency give better stability limit estimates than those of distant monitoring locations.

These points can be verified by the examining the results for Abitibi given on Table 5−6. Here, we compare the estimates obtained from the two signal energy formulas applied to results obtained at our usual monitoring location, Duvernay, and at the point of fault, Abitibi. Regardless of the formula used, the signal energy limit estimate gives better results when applied to a voltage waveform very close to the fault location. This may be due to the larger voltage swing experienced at locations near the contingency which translates into greater signal energy discrimination from one case to the next.

This observation also helps us understand the difficulty of obtaining a good signal limit estimate for faults at Churchill Falls: A complex network including long lines, and especially additional generation, sizeable load and active voltage support separates the distant Duvernay monitoring location from Churchill Falls, thereby degrading the signal energy readings at Duvernay for faults at Churchill Falls. In view of such obstacles, the 190 MW error of Table 5−4 must

TABLE 5–5

Comparison of Estimated and True[1] Stability Limits For Spontaneous Loss of Line

Contingency Location	*Points[2] (p.u.)*			*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*
Using 2–point formula (eq. 5–26):						
Albanel	.8188		.9547	8459	8830	4.20
Albanel		.9094	.9547	8836	8830	0.07
Using 3–point formula (eq. 5–23):						
Albanel	.8188	.9094	.9547	8769	8830	0.69

TABLE 5–6

Comparison of Estimated and True[1] Stability Limits: Effect of Monitoring Location (Fault With Loss of Line)

Fault Location	*Points[2] (p.u.)*			*Estimated Limit (MW)*	*True[1] Limit (MW)*	*|Error|[3] (%)*	*Monitoring Location*
Using 2–point formula (eq. 5–26):							
Abitibi	.9117		.9874	7867	7930	0.79	Duvernay
Abitibi	.9117		.9874	7883	7930	0.59	Abitibi
Using 3–point formula (eq. 5–23):							
Abitibi	.9117	.9622	.9874	7922	7930	0.10	Duvernay
Abitibi	.9117	.9622	.9874	7931	7930	0.01	Abitibi

1 *Limit obtained by simulation: within 25 MW*
2 *Points used to estimate limit, in p.u. of True Limit*
3 *% error with respect to stability limit obtained by simulation*

be seen as remarkably good. This proves the value of voltage–based signal energy limit estimation on the basis of few network monitoring points.

5.6.5 REGION OF VALIDITY OF SIGNAL ENERGY ANALYSIS

Figure 5–15 shows the results of Fig. 5–13 plotted on a per–unit basis with respect to the stability limit for each individual fault location. These show that signal energy exhibits a slightly sharper "cutoff" characteristic for faults accompanied by changes in topology, the sharp rise occurring within approximately 3.5% of the limit as opposed to 5% in Fig. 5–9. In addition, we have seen that the three–point formula gives good results provided that the farthest point is within 16% of the limit. These points appear to define the region of validity of the initial assumptions (i.e. dominant pole damping very near zero), as borne out by the results of Tables 5–3 and 5–4. From such observations, is it possible to give a more formal definition for the region of validity of the proposed signal energy model?

Figures 5–6, 5–9, 5–13 and 5–15 indicate that some measure of signal energy slope has the potential of establishing proximity to the stability limit. Let us examine this by first rewriting the simplified signal energy expression as follows (i.e. equation 5–26):

$$E_{ij}(P_i) = \frac{C_{0ij}}{C_{2ij}} \frac{-1}{(P_i/C_{2ij} - 1)} \qquad (5-36)$$

Since the power term in the denominator is now in per unit with respect to the stability limit,

$$P_i' = P_i/C_{2ij} \qquad (5-37)$$

equation (5–36) can be rewritten as:

$$E_{ij}(P_i') = \frac{C_{0ij}}{C_{2ij}} \frac{-1}{(P_i' - 1)} \qquad (5-38)$$

Let us now define per unit signal energy $E_{ij}'(P_i')$ as follows:

$$E_{ij}(P_i') = \frac{C_{0ij}}{C_{2ij}} E_{ij}'(P_i') \qquad (5-39)$$

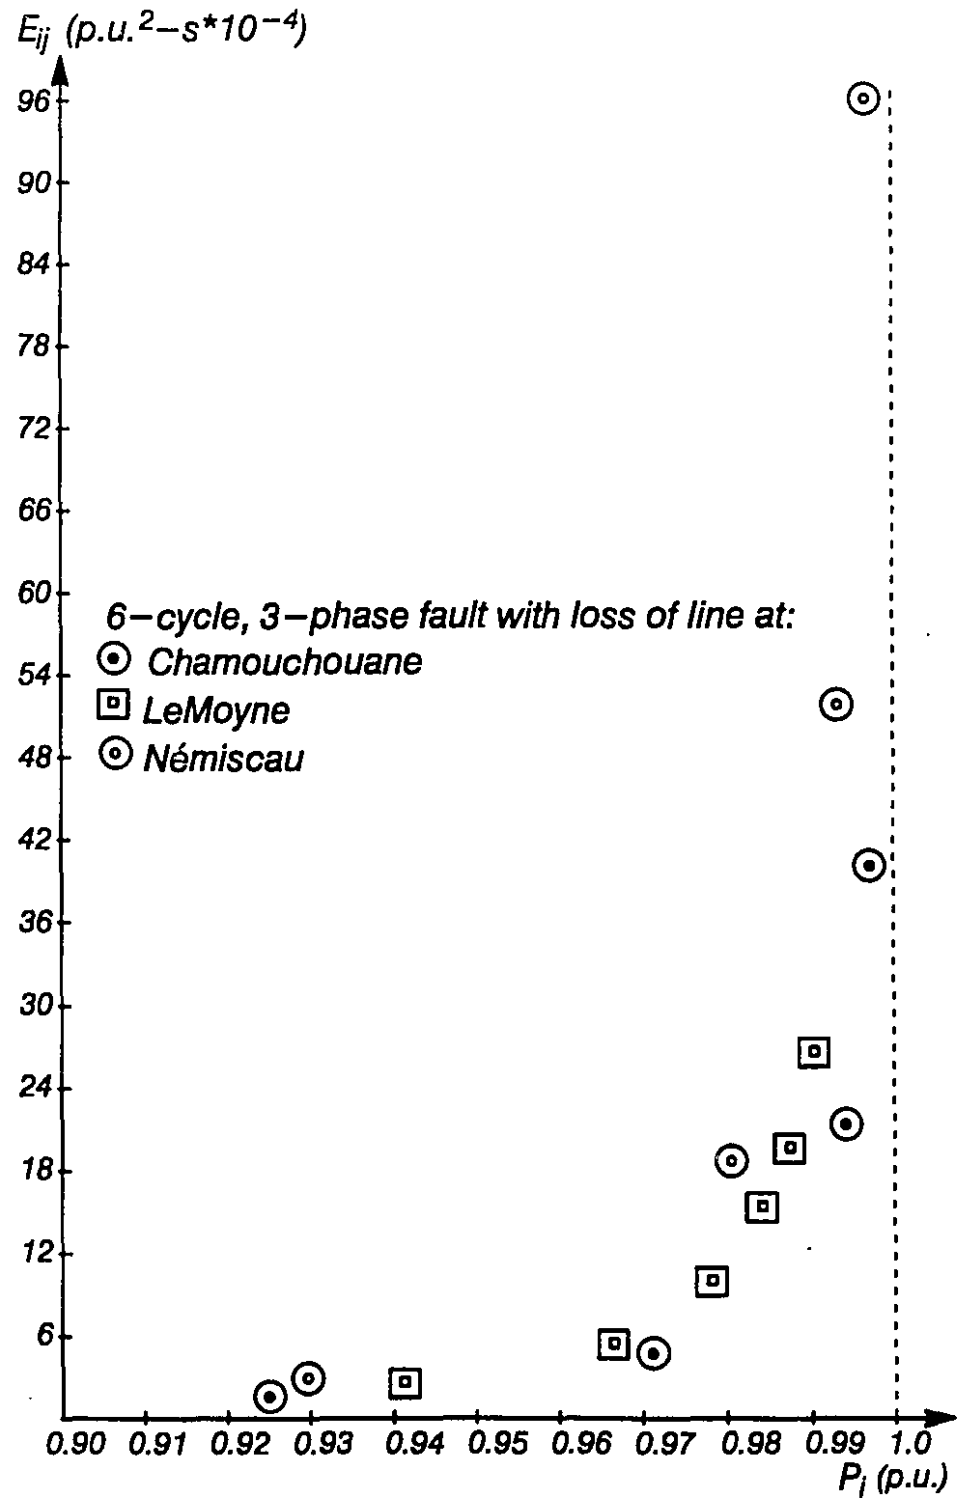

Fig. 5–15. Signal energy results of Fig. 5–13 plotted in per–unit of the simulated stability limit of each contingency (6–cycle, 3–phase fault with loss of line).

TABLE 5−7

Per Unit Signal Energy Slope

$$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'}$$

as a Function of Limit Estimate Error

Fault Location	*Points*[1] *(p.u.)*			*|Error|*[2] *(2−point formula) (%)*	*|Error|*[3] *(3−point formula) (%)*	$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'}$[4]
LG2	.8602	.9554	.9792	0.79	0.14	344
Le Moyne	.8920	.9414	.9661	1.42	0.84	273
Némiscau	.9293	.9807	.9936	0.26	0.05	2210
Albanel	.9293	.9550	.9807	0.99	0.64	733
Abitibi	.9117	.9622	.9874	0.79	0.10	899
Chibougamau	.9056	.9543	.9787	1.26	0.85	497
Chamouchouan	.9293	.9712	.9942	0.06	0.15	2438
La Vérendrye	.8382	.9217	.9635	4.49	0.86	169
Churchill Falls	.8976	.9472	.9721	2.61	2.36	350

1 *Points used to estimate limit (in p.u. of stability limit, seeTable 5−4)*
2 *% error in limit estimate obtained from 2−point formula (seeTable 5−3)*
3 *% error in limit estimate obtained from 3−point formula (seeTable 5−4)*
4 *Calculated using the closest and farthest points*

where

$$E_{ij}'(P_i') = \frac{-1}{(P_i' - 1)} \qquad (5-40)$$

Per unit signal energy is dependent only on per−unit power with respect to the stability limit. The accuracy of the signal energy estimate of Tables 5−3 and 5−4 can now be characterized in terms of per unit signal energy slope with respect to per unit power. Table 5−7 shows the value of per unit signal energy slope associated with the points used to estimate the limits in Tables 5−3 and 5−4 (i.e. for faults with loss of line). From this information, the minimum per unit signal energy slope which roughly defines the region of validity of signal energy analysis can now be determined. Using equation (5−40), provided that signal energy is monitored from a not−two−distant location (see 5.6.4), the

2−point formula (equation 5−26) gives limit estimates with less than 1.3% error when:

$$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'} \geq 497 \qquad (5-41)$$

For the same value of slope, the 3−point formula clearly gives less than 1% error in the limit estimate.

A final point. The per unit signal energy slope value given in equation (5−41) can also be assumed to mark the point at which signal energy begins its rapid rise, thereby formally circumscribing the region of validity of the proposed signal energy model. Taking the derivative of equation (5−40), per unit signal energy slope is given by:

$$\frac{dE_{ij}'(P_i')}{dP_i'} = \frac{1}{(P_i' - 1)^2} \qquad (5-42)$$

We now substitute the value for the slope given in (5−41) into equation (5−42) and solve for P_i', which gives:

$$P_i' = 0.955 \text{ p.u.} \qquad (5-43)$$

The knee of the signal energy function can therefore be taken to begin at approximately 95.5% of the stability limit.

5.6.5.1 Rapid estimates of transient stability transfer limits

The preceding results now provide all the elements permitting the determination of highly−accurate transient stability transfer limit estimates with at most three stable simulations (i.e. stable simulations take the longest time to execute, see 6.3.2.2).

The stategy is as follows:

1) One prepares a base case using a typical operating point used in operations planning studies, such as peak load conditions (see 2.4.1.3).
2) Having obtained two stable simulations, one calculates the associated

signal energy of each one.

3) Using the signal energy data thus obtained and the 2−point signal energy formula of equation (5−26), one obtains an estimate of the limit, C_{2ij}.
4) Using equation (5−37), one transforms the values for power in p.u. of the estimated stability limit.
5) One calculates the p.u. signal energy of each point using (5−40) and, thereafter, the p.u. signal energy slope with respect to p.u. power:

$$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'}$$

6) One then applies the test proposed in equation (5−41): If the calculated slope is greater than the indicated value, the two points give an estimate with less than 1.3% error. If the slope is less than the indicated value, an additional simulation is required in the neighbourhood of the estimated limit value. One then returns to 2) above and repeats the process using the two points closest to the limit. Of course, when three points are available, the 3−point formula can be used, ensuring even higher accuracy in the limit estimate.

This is essentially the process outlined in the **accept** filter example of Fig. 3−16. Because such a process requires at most three stable simulations and yields a low−error stability limit, it has the potential of achieving rapid estimates of power transfer limits in the on−line environment

5.6.6 STABILITY MARGIN

Figure 5−15 again illustrates how the concept of stability margin works. From these results, one sees that the network is many times more fragile at 99% than at 96.5% of the limit.

5.6.7 NON−LINEAR ACTIVE VOLTAGE SUPPORT

Figure 5−15 also exhibits differences in curvature among the various curves, as in Fig. 5−9. The three fault locations plotted here include one equipped with 2 SVCs (Chamouchouane), one very near a major generating station (LG 4) equipped with static excitation systems (Le Moyne) and one having neither SVCs nor SCs (Némiscau). As discussed in 5.5.4, the signal

energy cutoff characteristic for fault locations equipped with SVCs or static excitation systems appears to be sharper than for locations with no such equipment. Translated in other terms, non–linear active voltage support tends to distort p.u. signal energy somewhat with respect to the model of equation (5–40).

5.7 SUMMARY

This chapter shows how new–generation tools such as frameworks can help perform research to improve our understanding of fundamental power systems phenomena.

First, it is shown that the network transient response to a 6–cycle, 3–phase fault is in fact a good approximation of its impulse response, apart from a linear scaling factor. Then, using the ELISA prototype shell as a laboratory tool, it is shown that the signal energy of the network impulse response E acts as a barometer io define the relative severity of a contingency with respect to some parameter P, for instance power generation or power transfer. For a given contingency, as the parameter is varied and a network approaches instability, signal energy increases smoothly and predictably towards an asymptote which defines a network's stability limit: This limit, in turn, permits us to compare the severity of different contingencies.

Using a Fourier transform approach, a simple function is derived which estimates network stability limits with surprising accuracy from only two or three sample simulations, provided that at least one of these is within a few percent of the stability limit. It is found that these results hold for all normal contingencies (i.e. faults with no change in topology, faults accompanied by subsequent loss of line or spontaneous loss of line), notwithstanding the presence of many active, nonlinear voltage–support elements (i.e. generators, synchronous condensers, SVCs, static excitation systems, etc.). Additionally, it appears that signal energy behaviour provides a powerful model for understanding the effectiveness of such long–used concepts as stability margin, and it may be possible to quantify the impact of different voltage support technologies in terms of the curvature of the signal energy cutoff characteristic.

A per unit signal energy function was also defined whose slope can be used

to determine proximity to the stability limit and which circumscribes the region of validity of the signal energy model. In particular, on the basis of per unit signal energy slope, the knee of the signal energy curve was found to begin at approximately 95.5% of the stability limit. Using these results, a strategy was proposed permitting the estimation of highly accurate transient stability transfer limits from at most three stable simulations.

The richness in new concepts that have arisen as a direct result of this research proves the effectiveness and importance of building frameworks capable of performing many hundreds of simulations which normally would have taken years. Of course, this research will eventually be fed back directly into the ELISA prototype in order to enhance existing limit–search processes.

CHAPTER 6

ON EVALUATING POWER SYSTEM STABILITY

6.1 INTRODUCTION

Though transient stability transfer limit determination has been an important focus of this thesis, the corollary issue of evaluating stability from time–domain simulation waveforms has only been treated indirectly. As we have seen, limits are primarily determined on the basis of acceptability, meaning that the post–contingency voltage (or frequency) at various network monitoring locations must be within traditional emergency operating limits a short time after the occurrence of the contingency.

However, as we have seen in the previous chapter, it is of theoretical and practical interest to determine the stability limit of a transmission system. In the interest of optimizing power transfer, rather than determining stability limits on the basis of long–established voltage– or frequency–dependent acceptability criteria, it may be preferable to 1) find the network stability limit and 2) impose a security margin based on some quantified measure of system robustness such as per unit signal energy slope. The key to mechanizing such a process resides in a criterion capable of detecting instability.

In this chapter, a frequency–domain stability criterion is proposed, based on the fact that a network's behaviour near the stability limit is governed by the action of its dominant poles, as observed in the previous chapter. Essentially, the criterion requires monitoring the Fourier transform of the network's transient response $R(\omega, P)$: When generation P is varied and the network crosses the stability limit, the angle of $R(\omega, P)$ fundamentally changes its behaviour, passing from a clockwise to a counterclockwise rotational behaviour. This is confirmed by results obtained from performing stability limit searches on the Hydro–Québec system using ELISA. Once again, we shall see that these results hold for normal contingencies, notwithstanding the presence of many active, nonlinear voltage support elements in the network.

6.2 FREQUENCY DOMAIN ANALYSIS OF NETWORK TRANSIENT RESPONSE NEAR THE STABILITY LIMIT

In the previous chapter, signal energy behaviour near the stability limit was explained effectively in terms of the two–dominant–pole transient response model of equation (5–16):

$$R_{ij}(\omega) = \frac{K_{ij}}{(j\omega - p_{ij})\,(j\omega - p_{ij}^*)} \qquad (6-1)$$

where p_{ij} and p_{ij}^* represent the dominant poles of the network in response to a normal contingency applied at location i and monitored at location j. From a linear systems standpoint, a stable system corresponds to the situation where the poles of equation (6–1) are on the LHP:

$$p_{ij} = \sigma_{ij} + j\omega_{ij}, \quad \sigma_{ij} < 0,\ \omega_{ij} > 0 \qquad (6-2)$$

Assuming, as in Chapter 5, that

$$\sigma_{ij} = \sigma_{ij}(P_i) \qquad (6-3)$$

an increase in generation P_i causes damping σ_{ij} to increase positively until the latter intersects the $j\omega$ axis at $\sigma_{ij} = 0$ and

$$P_i = C_{2ij} \qquad (6-4)$$

As before, C_{2ij} is the stability limit. If power is increased beyond the stability limit, p_{ij} and p_{ij}^* cross the $j\omega$ axis, causing the pole damping components to take on positive values:

$$\sigma_{ij} \geq 0,\ P_i \geq C_{2ij} \qquad (6-5)$$

Let us now consider the anticipated behaviour of the polar plot of equation (6–1) as angular frequency is varied over the range:

$$0 \leq \omega < \infty \qquad (6-6)$$

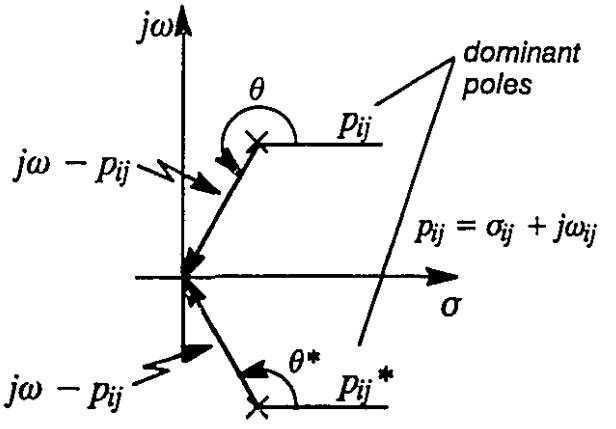

(a) $\omega = 0 : \theta + \theta^* = 360^0 \Rightarrow \angle R = -360^0$

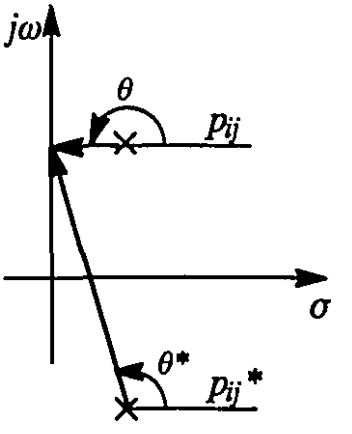

(b) $\omega = \omega_{ij} : 270^0 < \theta + \theta^* < 360^0 \Rightarrow -360^0 < \angle R < -270^0$

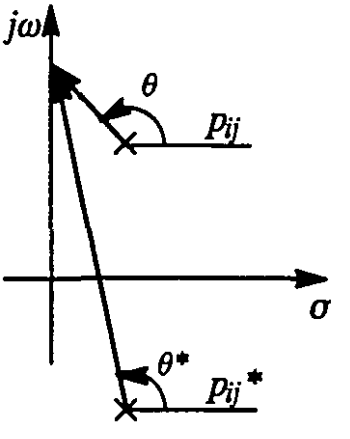

(c) $\omega > \omega_{ij} : \theta + \theta^* \rightarrow 180^0 \Rightarrow \angle R \rightarrow -180^0$

Fig. 6–1. s–plane plot showing how dominant RHP poles near the $j\omega$ axis contribute to $R_{ij}(\omega)$ phase as a function of angular frequency ω. Distant poles and zeroes impact the polar plot (Fig. 6–2) at $\omega = 0$ (i.e. initial phase) and for values $\omega >> \omega_{ij}$.

and as the poles p_{ij} and p_{ij}* take on different positions in the s−plane.

6.2.1 ANTICIPATED POLAR PLOT BEHAVIOUR OF DOMINANT POLE MODEL

Figure 6−1 shows the dominant poles positioned on the RHP of the pole−zero plot, just after having crossed the imaginary axis. Let us now consider the phase behaviour of $R_{ij}(\omega)$ as we vary angular frequency along the positive $j\omega$ axis (Fig. 6−2). In this analysis, we implicitly assume that the dominant poles are so close to the $j\omega$ axis that all other network poles and zeroes appear distant and consequently do not contribute to $R_{ij}(\omega)$ phase changes. For simplicity, let us also assume that, at zero frequency, the phase contribution of all other poles and zeroes is zero.

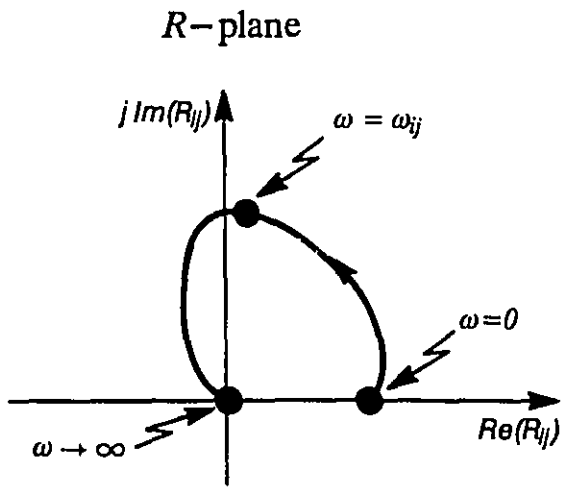

Fig. 6−2. Sketch of polar plot for two−dominant−pole output spectrum $R_{ij}(\omega)$ as ω is varied as shown in Fig. 6−1. The basic characteristic of this plot is counterclockwise (CCW) phase rotation.

From Fig. 6−1 (a), at 0 radians/second, the combined contribution of the dominant poles to $R_{ij}(\omega)$ phase is -360^0. Examining Fig. 6−1 (b), the contribution of the dominant poles to $R_{ij}(\omega)$ phase is highest in the region between 0 and the dominant pole natural frequency, ω_{ij}: In this band of frequencies, the pole angles change rapidly causing $R_{ij}(\omega)$ phase to increase to some value between -270^0 and -360^0, resulting in a counterclockwise (CCW) departure on the R−plane plot (Fig. 6−2). When higher frequencies are considered, Fig. 6−1 (c) leads us to expect that $R_{ij}(\omega)$ phase continues to progress counterclockwise,

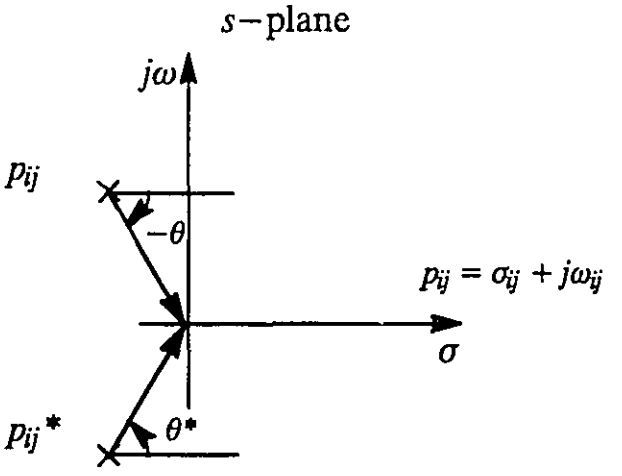

(a) *s−plane plot showing effect of LHP dominant poles near jω axis: at ω = 0 , θ + θ* = 0 0 => ∠R = 0 0*

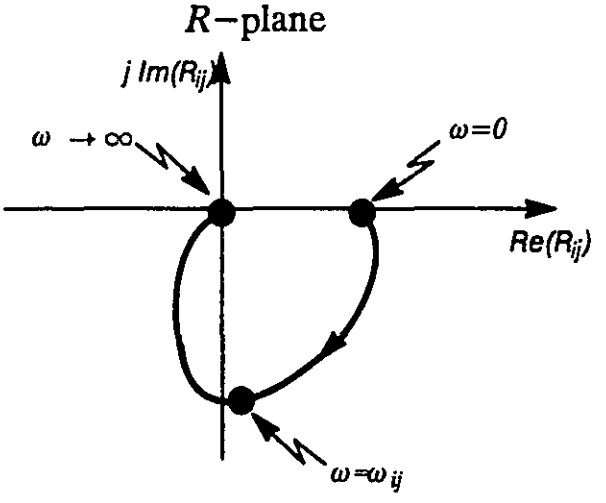

(b) *Polar plot of $R_{ij}(\omega)$ for LHP dominant poles near jω axis.*

Fig. 6−3. Sketch of polar plot to illustrate the principle of how dominant LHP poles near the $j\omega$ axis contribute to $R_{ij}(\omega)$ phase. The basic characteristic of this plot is clockwise (CW) phase rotation.

tending towards a -180^0 phase angle, resulting in a maximum possible -90^0 contribution per RHP pole as sketched in Fig. 6–2. A corollary to this is that for two dominant LHP poles very close to the imaginary axis, $R_{ij}(\omega)$ phase should rotate clockwise (CW) on the R–plane (Fig. 6–3).

Clearly, for network transfer functions with RHP poles close to the imaginary axis, counterclockwise rotation of $R_{ij}(\omega)$ phase on the R–plane about the origin, at least up to the dominant–pole natural frequency, appears to be the characteristic behaviour. Though we have illustrated this property for two–pole output spectra, one can arrive at similar conclusions for multiple RHP poles close to the $j\omega$ axis.

To obtain $R_{ij}(\omega)$ for the purpose of observing this behaviour, one need only perform a Fourier transform of the network transient response as defined by equation (5–27). One then plots the imaginary versus the real parts of $R_{ij}(\omega)$ as illustrated in Fig. 6–2 to obtain the polar plot.

6.2.2 SECONDARY POLE–ZERO INTERACTIONS

As has already been pointed out, real networks have far more than two poles, a large number of zeroes, and evidence of this can be seen in the various frequency spectra found in Chapter 5. This suggests that the previous analysis must be adjusted somewhat to anticipate the impact of additional poles and zeroes near dominant poles.

In the case of stable networks, multiple LHP poles should be evidenced by CW rotation of the polar plot about the origin, amounting to 90^0 per LHP pole.

In the case of unstable networks, multiple RHP poles should be evidenced by CCW rotation of the polar plot about the origin, amounting to 90^0 per RHP pole. If LHP poles are sufficiently near the $j\omega$ axis for their effect to be observable, dominant RHP poles will exhibit initial CCW rotation, followed by a CW 90^0 contribution per LHP pole, as sketched in the example of Fig. 6–4. Of course, such sketches are highly dependent on relative pole position.

LHP or RHP zeroes may occasionally lie sufficiently close to the $j\omega$ axis as to impact the polar plot. Complex zeroes, like poles, occur in conjugate pairs

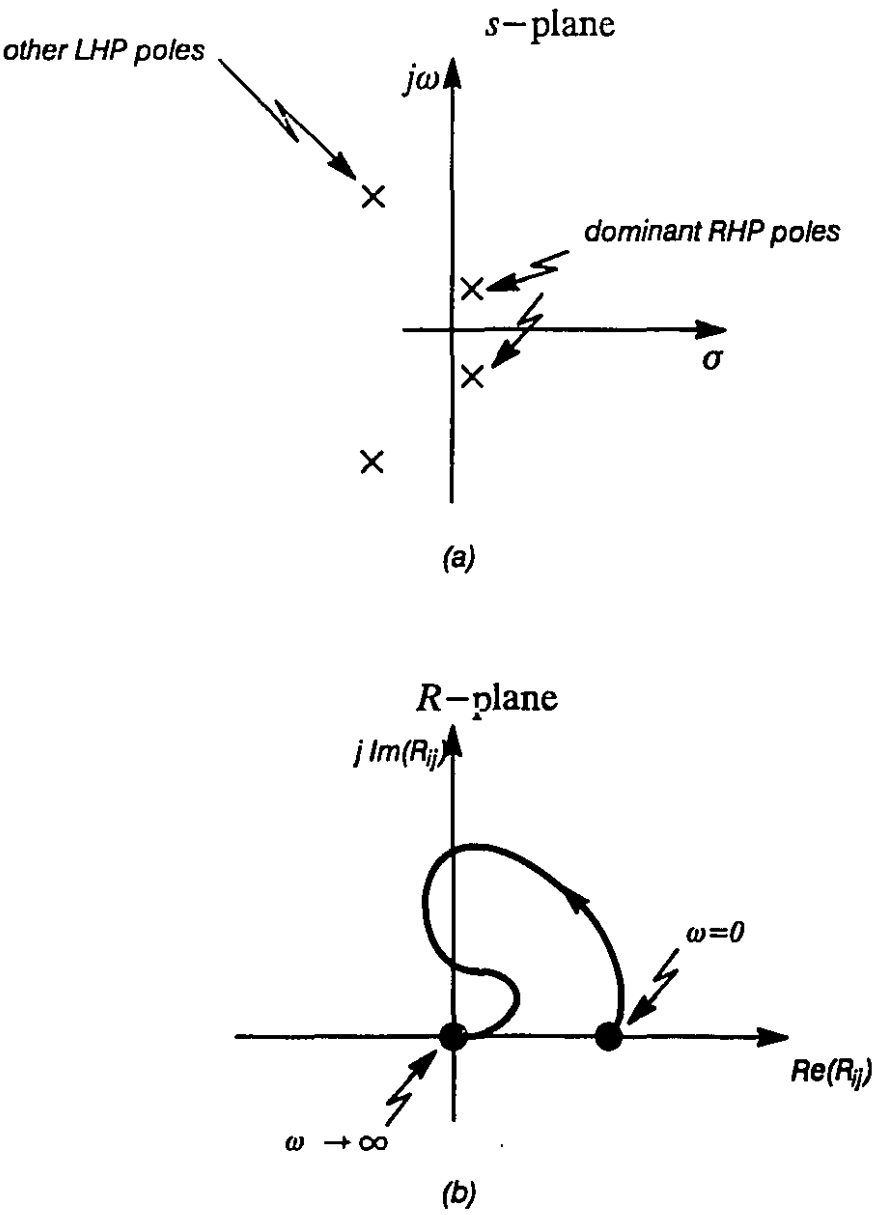

Fig. 6−4. Sketch of polar plot illustrating the principle of dominant RHP poles near the $j\omega$ axis interacting with other LHP poles. Note the initial CCW rotation due to the dominant RHP poles, followed by subsequent CW rotation.

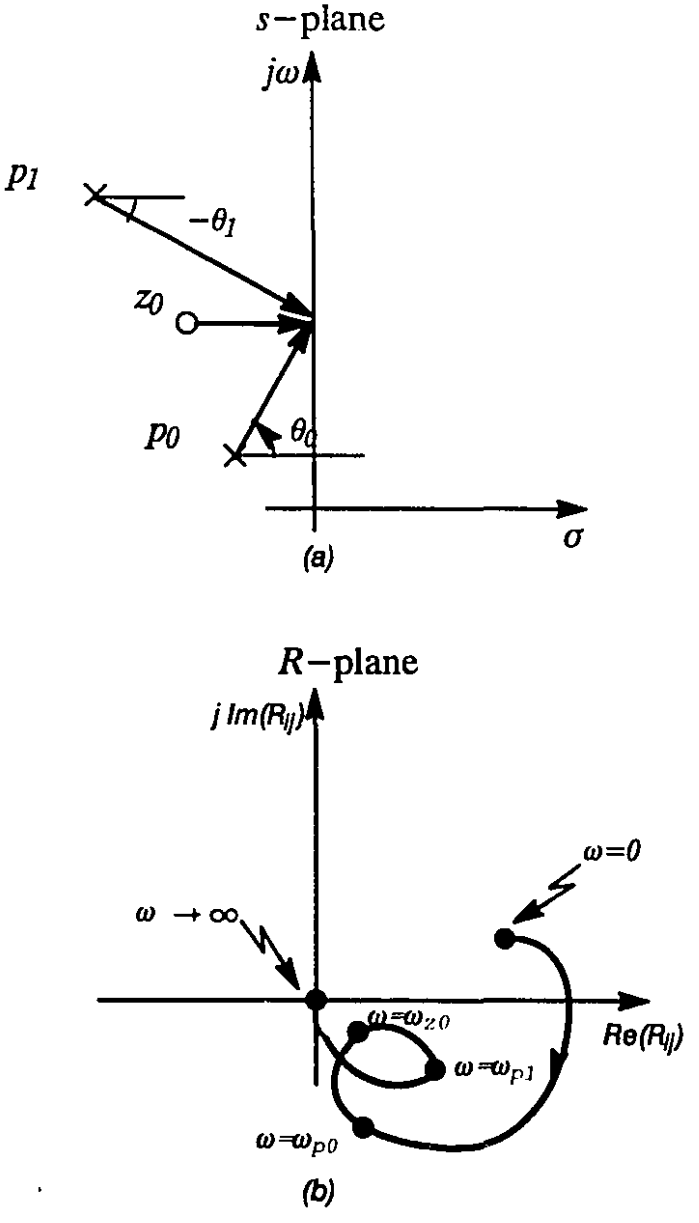

Fig. 6–5. Sketch of polar plot illustrating the impact of a zero between two poles in the LHP: A local CW rotation results, displaced from the origin.

as a consequence of their originating from rational algebraic functions [Saucedo & Schiring 1968]. As a result, zeroes will contribute either a 90^0 CCW rotation per LHP zero, or a 90^0 CW rotation per RHP zero to the polar plot. In the case of dominant LHP poles, the CCW rotational contribution of LHP zeroes will be evidenced as small CW rotations displaced from the origin and superimposed on the plot due to decreasing $|R_{ij}(\omega)|$ as ω passes near the zero, as illustrated in Fig. 6–5. RHP zeroes also cause a decrease in $|R_{ij}(\omega)|$ as ω passes near the zero, but the CW angle of rotation remains the same. A similar analysis can be performed involving RHP and LHP zeroes interacting with dominant RHP poles except that the dominant sense of rotation is CCW.

6.2.3 VALIDITY OF THE APPROACH

As a general rule, RHP poles contribute to CCW phase rotation in the R–plane, but LHP zeroes also have the same effect. In addition, both LHP poles and RHP zeroes contribute to CW rotation. Notwithstanding the previous analysis, polar plot interpretation seems to lend itself to some confusion. This being the case, how can the previous analysis be of any significance?

Consider the effect of parameter P_i on a possible migration of zeroes in the s–plane. If zeroes were significant, they would impact signal energy as the network approaches the limit. However, near the stability limit, a two–pole transfer function, ignoring zeroes, has been shown to explain signal energy behaviour with high accuracy. This does not mean that zeroes do not move under the influence of P_i. It only means that, near the limit, dominant poles overide all other contributions. Consequently, if time–domain simulations are sufficiently near the stability limit that the previous chapter's signal energy model is applicable, then $R_{ij}(\omega)$ phase angle analysis can also be taken to be valid.

In the previous chapter, we saw that the region of validity of signal energy analysis (i.e. equation 5–23) could be given in terms of a minimum value of per unit signal energy slope:

$$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'} \geq 497 \qquad (5\text{–}41)$$

This criterion corresponds to all values of power greater than 0.955 p.u. (i.e.

in per unit of the stability limit). For the moment, this can be taken to define the region of validity of the above polar plot analysis.

6.3 A CASE STUDY: THE HYDRO–QUÉBEC NETWORK

6.3.1 TIME– AND FREQUENCY–DOMAIN RESULTS

Let us consider the time–domain network response obtained in the course of the stability limit search of Figs. 5–1 (a) and (b) for faults with no loss of line, and in Figs. 5–14 (a) and (b) for faults with loss of line. The network response, as before, is given by equation (5–27):

$$r(t, P) = v(t, P) - V_{st-st}(P) \qquad (5-27)$$

These time–domain waveforms are plotted with points at 2 cycle intervals (i.e. Δt = 0.03333 seconds) and fast Fourier transforms are performed after the fault has been cleared (i.e. after 6 cycles). Though 600 cycle simulations have been found adequate for stability determination on the Hydro–Québec system, all stable polar plots were obtained from 3000 cycle simulations. A comparison of 600 and 3000 cycle polar plots appears in Fig. 6–6 for the same stable case: The longer simulation time is seen to generate a smoother, less discontinuous plot; nevertheless, CW phase rotation is clearly evident in both cases.

Figures 6–7 to 6–12 illustrate typical results. In particular, Figs. 6–7 and 6–8 show polar plots of the fast Fourier transform of the transient response for some of the key simulations of Figs. 5–1 (a) and (b) (i.e. 3–phase fault with no change in topology). Figures 6–9 and 6–10 show the same for the simulations of Figs. 5–14 (a) and (b) (i.e. 3–phase fault with loss of line). Figure 6–11 shows the polar plots for the case of a spontaneous loss of line at Albanel and Fig. 6–12 presents the corresponding transient stability simulation waveforms.

A final point. Fast Fourier transforms were obtained using the MATLAB software [Little & Shure 1992]. The true Fourier transform (*FT*) differs from the MATLAB–computed fast Fourier transform (*FFT*) by a constant, Δt:

$$FT = FFT\ \Delta t \qquad (6-7)$$

where, in our case,

$$\Delta t = .03333 \text{ seconds} \qquad (6-8)$$

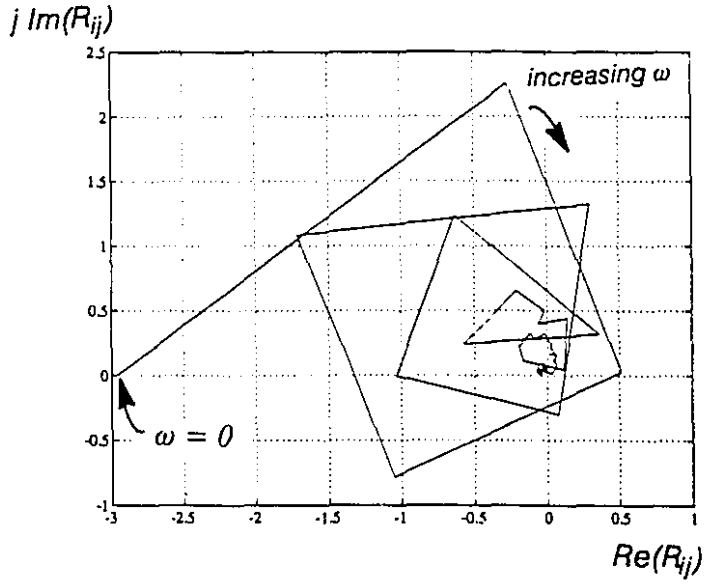

(a) Typical polar plot obtained from 600 cycle simulation

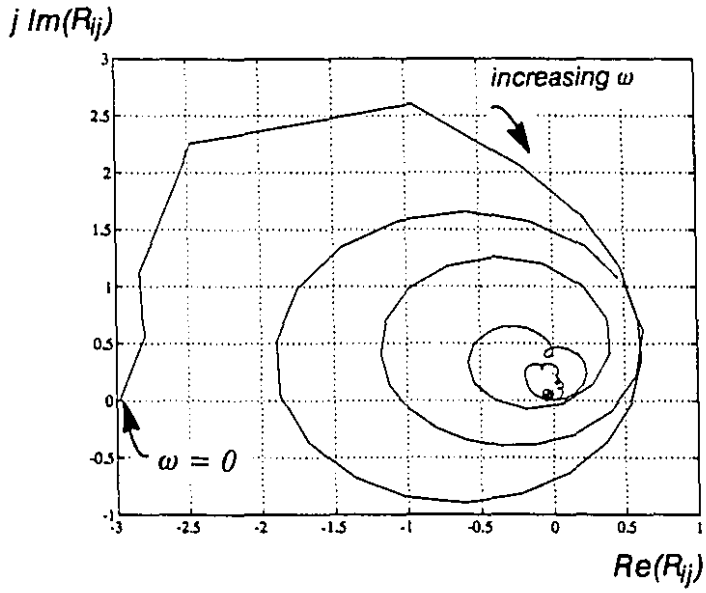

(b) Typical polar plot obtained from 3000 cycle simulation

Fig. 6−6. Two examples illustrating the sensitivity of stable *FFT* polar plots with respect to simulation time length. CW rotation is preserved on the simulation of shorter duration, though the plot is less continuous. These plots correspond to the base case + 350 MW for a 6−cycle, 3−phase fault at Abitibi with no loss of line. Voltage is monitored at Duvernay (i.e. stability limit of Fig. 5−1).

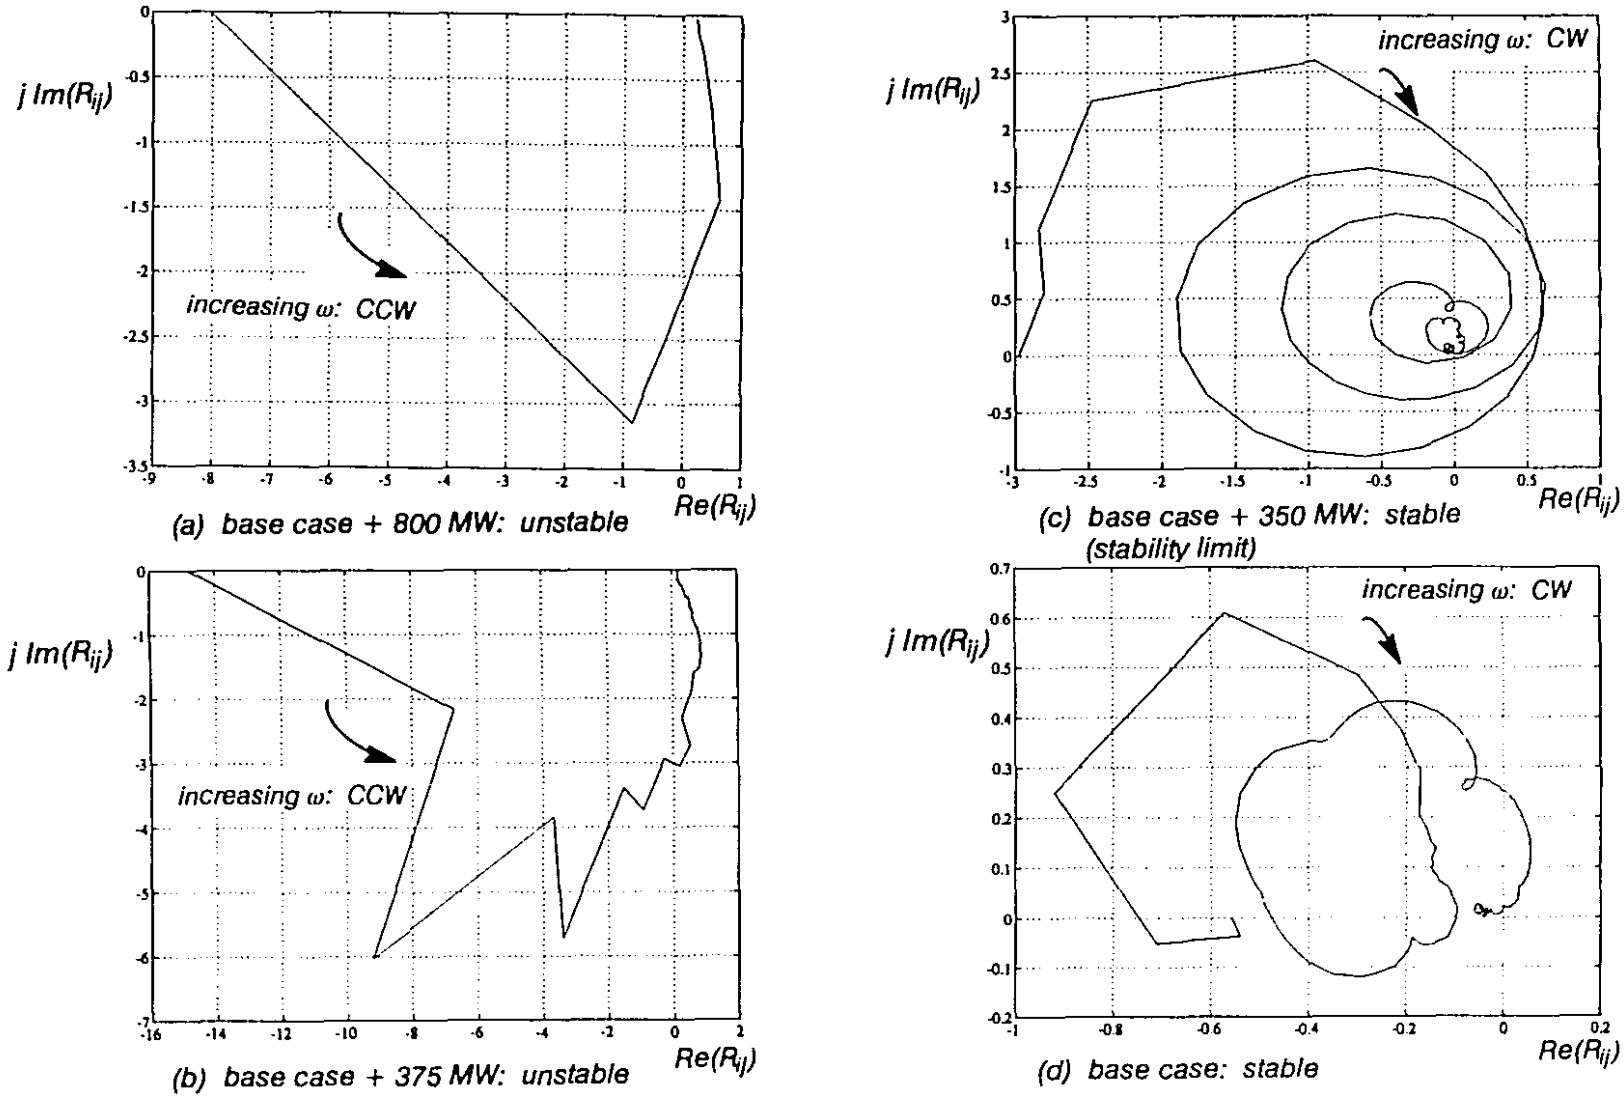

Fig. 6–7. *FFT* polar plots for key Duvernay voltage waveforms of Fig. 5–1 (a) for 6–cycle, 3–phase fault at Abitibi and no loss of line. Unstable cases exhibit CCW phase rotation; stable cases show CW rotation.

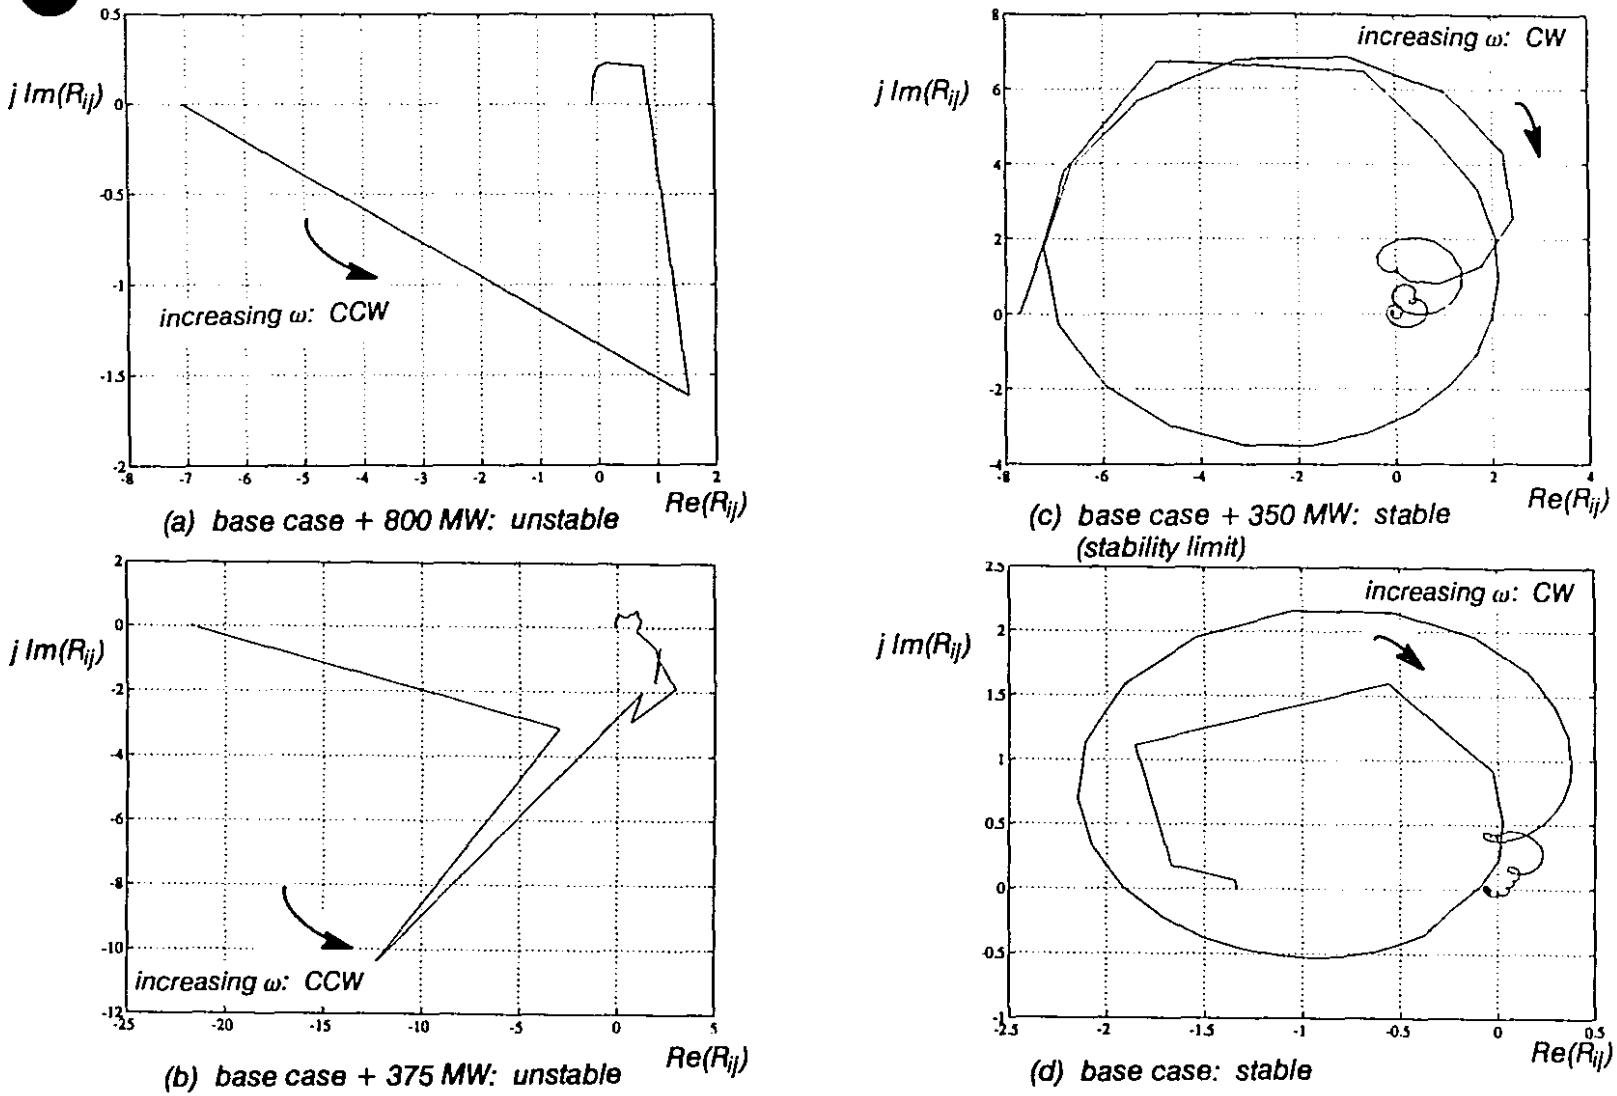

Fig. 6−8. *FFT* polar plots for key Abitibi voltage waveforms of Fig. 5−1 (b) for 6−cycle, 3−phase fault at Abitibi and no loss of line. Unstable cases exhibit CCW phase rotation; stable cases show CW rotation.

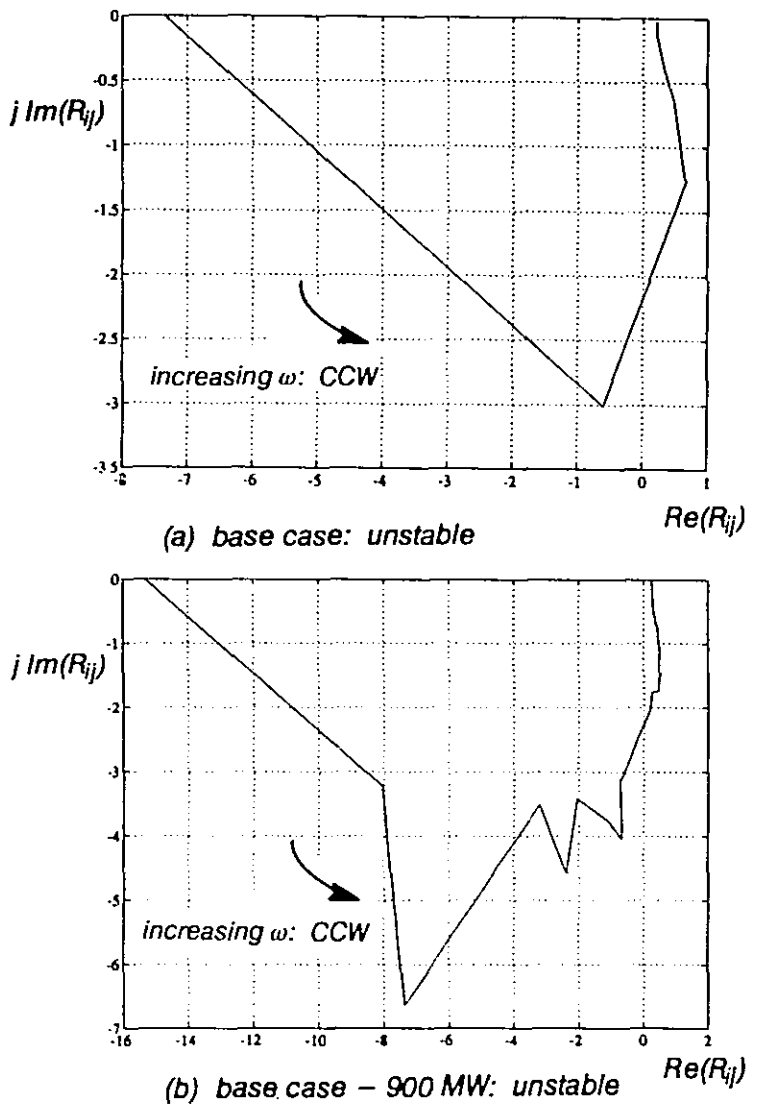

Fig. 6–9. *FFT* polar plots for key Duvernay voltage waveforms of Fig. 5–14 (a) for 6–cycle, 3–phase fault at Abitibi with loss of line. Unstable cases exhibit CCW phase rotation; stable cases show CW rotation.

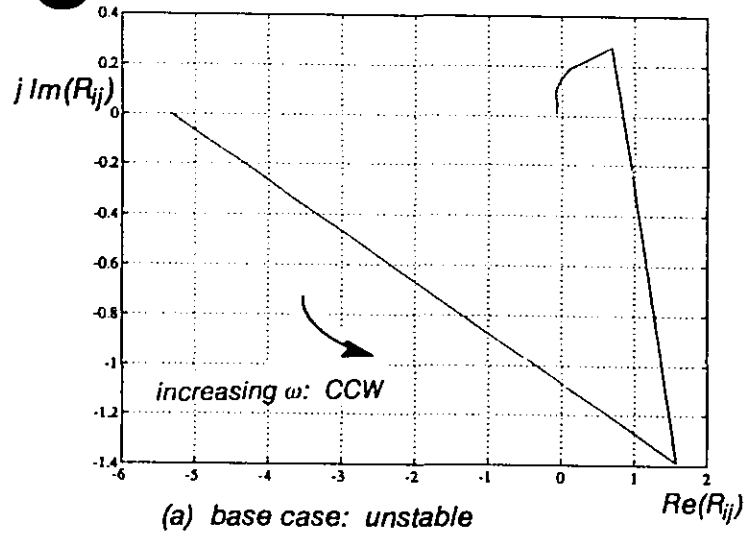

(a) base case: unstable

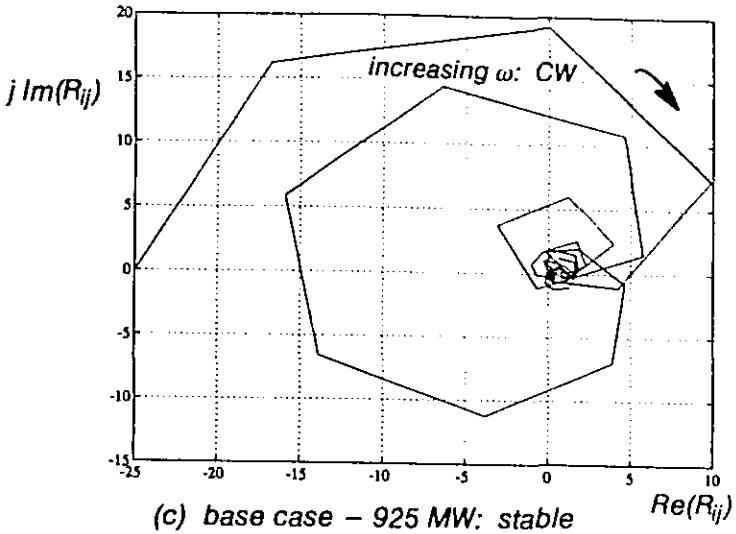

(c) base case – 925 MW: stable (stability limit)

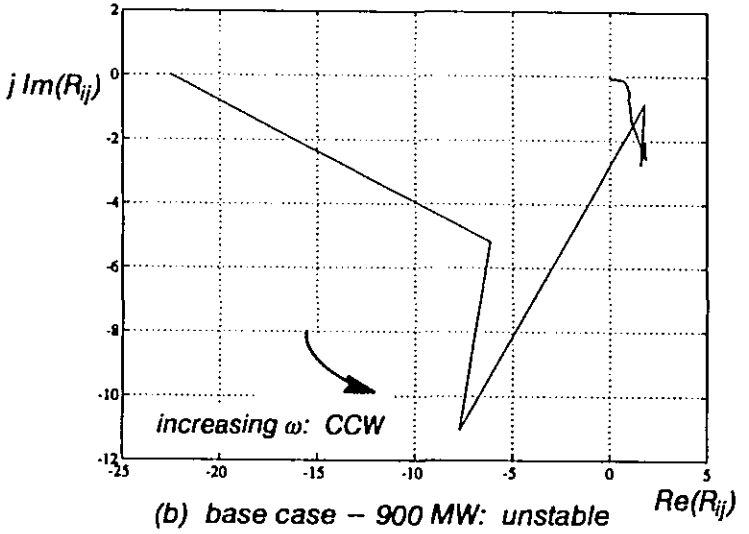

(b) base case – 900 MW: unstable

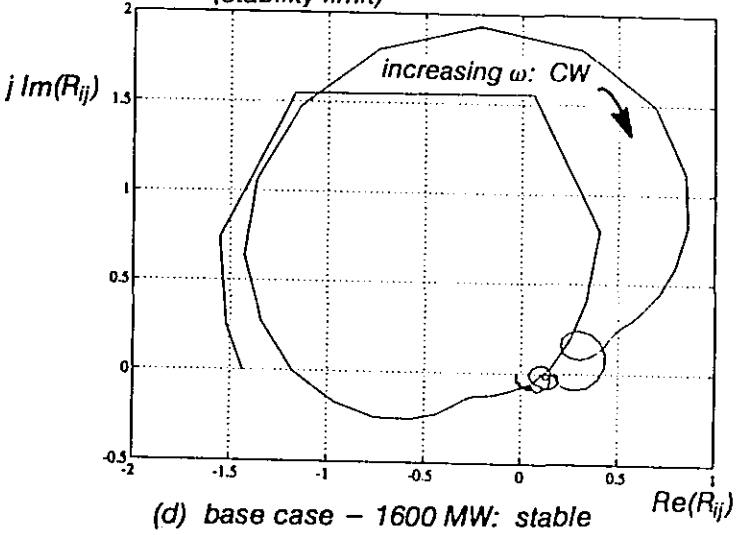

(d) base case – 1600 MW: stable

Fig. 6–10. *FFT* polar plots for key Abitibi voltage waveforms of Fig. 5–14 (b) for 6–cycle, 3–phase fault at Abitibi with loss of line. Unstable cases exhibit CCW phase rotation; stable cases show CW rotation.

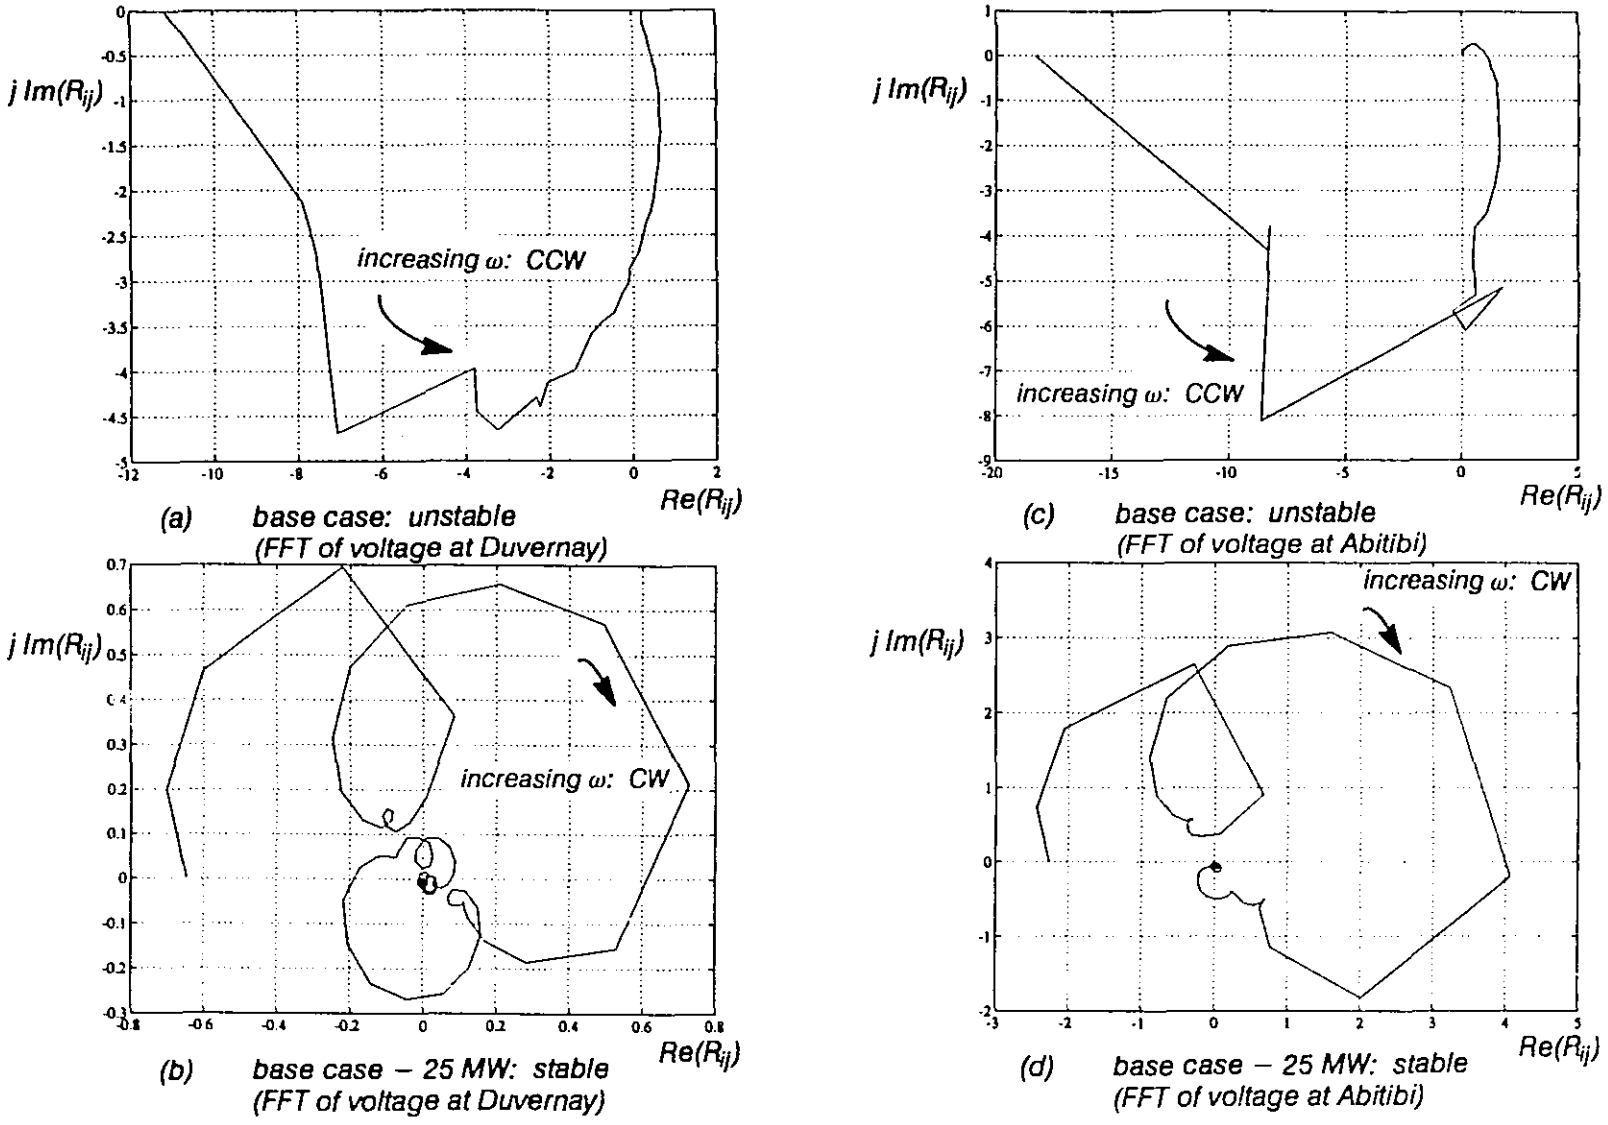

Fig. 6–11. *FFT* polar plots for key time–domain voltage waveforms of Fig. 6–12 for spontaneous loss of line at Albanel. Unstable cases exhibit CCW phase rotation; stable cases show CW rotation.

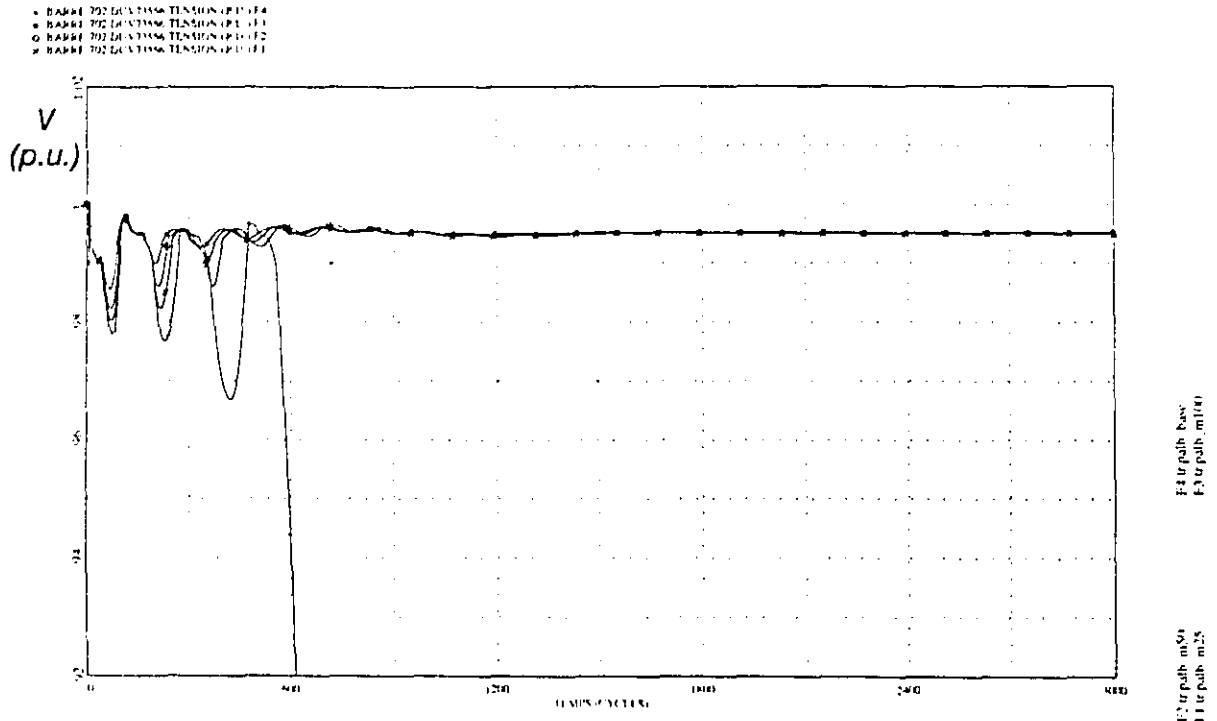

(a) Voltage monitored at Duvernay

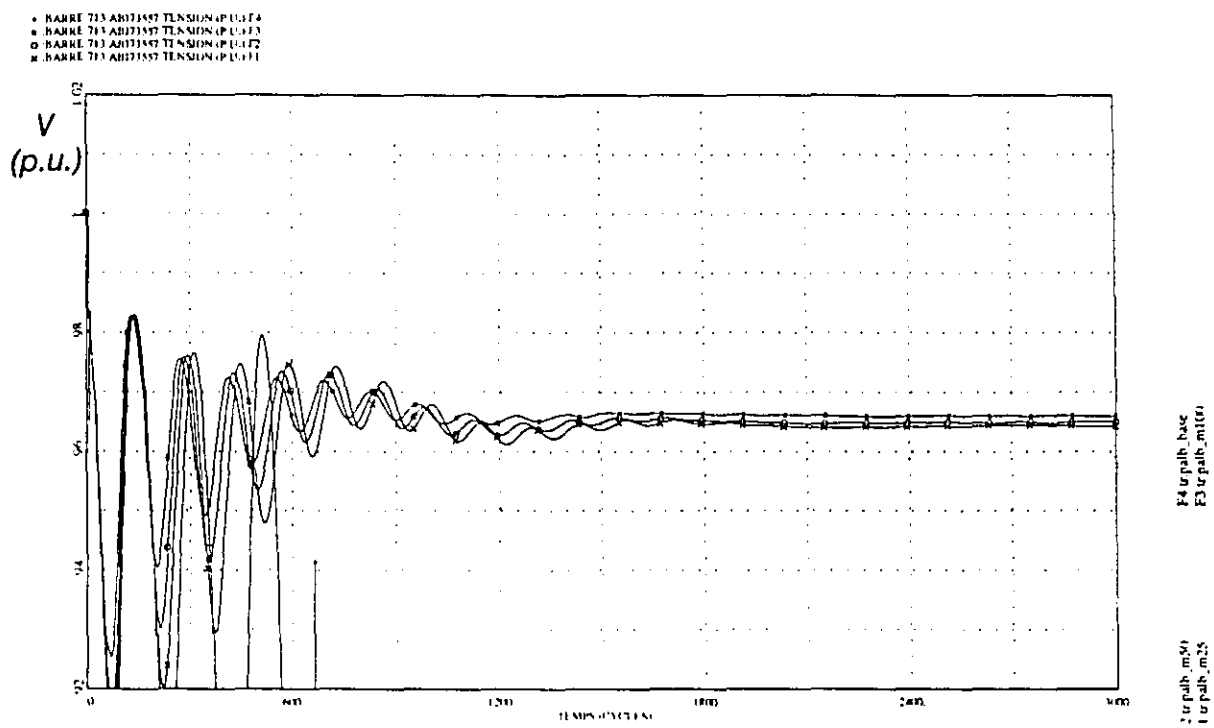

(b) Voltage monitored at Abitibi

Fig. 6–12. Transient stability transfer limit search performed by ELISA for spontaneous loss of line at Albanel (and no fault).

6.3.2 DISCUSSION OF RESULTS

6.3.2.1 STABLE CASES

As previously suggested, all stable cases exhibit CW rotational phase angle behaviour. This is true whether the case is far from the stability limit (see Figs. 6–7 (d), 6–8 (d), 6–9 (d) and 6–10 (d)) or within 25 MW of the limit (see Figs. 6–7 (c), 6–8 (c), 6–9 (c) and 6–10 (c)), or whether voltage is monitored at different locations (see Fig. 6–7 versus 6–8, and Fig. 6–9 versus 6–10). Multiple poles are evident from the inward CW spiralling behaviour of the plots, though the initial 90^0 contributions – due to dominant poles – exhibit the greatest amplitude. LHP zero contributions are also evident from the existence of localized CW rotations, displaced from the origin. From the point of view of our previous linear–based analysis, stable cases harbour no surprises.

6.3.2.2 UNSTABLE CASES

Contrary to the sketches of Figs. 6–2 and 6–4, the first point on the plots of unstable cases begins at a phase angle of -180^0 rather than at -360^0. This is because the zero–frequency (i.e. DC) component of the time–domain response is negative with respect to the post–contingency steady state voltage.

In general, unstable cases exhibit CCW polar plot rotational phase behaviour, not only initially, but through all frequencies, typically falling into the origin (i.e. for increasing frequency) with little or no contribution from other poles and zeroes. This observation is borne out for cases at some distance from the stability limit (see Figs. 6–7 (a), 6–8 (a), 6–9 (a) and 6–10 (a)) or within a few MW of the limit (see Figs. 6–7 (b), 6–8 (b), 6–9 (b) and 6–10 (b)). This also appears to be the case for faults with no loss of line (Figs. 6–7 and 6–8), faults with loss of line (Figs. 6–9 and 6–10) and spontaneous loss of line (Fig. 6–11).

Because unstable cases are interrupted prematurely by the ST600 transient stability program (i.e. whenever two generators are separated by more than some user–specified angle, the ST600 stability software ceases its execution, thereby prematurely interrupting the simulation), as is the case for many other commercial transient stability software, the time interval over which the Fourier transform is determined is generally shorter than the requested simulation time, and this is reflected in the polar plots in that they are frequently characterized

by discontinuous jumps, particularly in the lower frequency region (i.e. the first few points). For example, it is difficult to ascertain whether the localized rotation in Fig. 6–10 (b) is due to the effect of some LHP pole or an insufficiently long time domain simulation. Extremely short simulation intervals are characteristic of very unstable cases: However, if the voltage waveforms at some monitoring location clearly exhibit transient voltage collapse as in Fig. 5–1 (a) (i.e. base case + 800 MW), the polar plots exhibit an equally clear CCW rotation as in Fig. 6–7 (a).

Because of the reduced simulation time length of unstable cases, the voltage at certain monitoring locations can occasionally give ambiguous, erratic or even erroneous time– and frequency–domain information. For example, the voltage at points very near large generating stations may be better controlled due to the effect of very powerful and rapid static excitation systems, causing the time–domain waveforms, though distorted, to appear stable in a truncated simulation. In such cases, generator frequency, or the voltage at intermediate locations between large power stations (see 2.4.9), are preferable to obtain a clear indication of instability in the time and frequency domain. One therefore concludes that, for voltage to be effective in stability evaluation, it must be monitored at a number of locations, preferably at intermediate points between large generating stations, to counter the effect of reduced simulation time–length.

A final remark. The CCW rotational behaviour of unstable cases when monitored at appropriate locations is characterized by the near–absence of secondary pole–zero interactions. When the dominant poles cross the $j\omega$ axis, system instability appears to be such an overwhelming phenomenon that contributions from other poles and zeroes become difficult to detect (see Figs. 6–9 (a) and (b) versus 6–10 (a) and (b)).

6.3.2.3 POLES AND ZEROES AT THE STABILITY LIMIT

Beyond the stability limit, one must be very careful with polar plot interpretation, particularly in relation to secondary pole–zero interaction. The very mechanism by means of which a power system enters instability – loss of synchronism – defines the limit of applicability of transient stability simulation methods. This is easily understood: State–of–the–art transient stability

software assumes that the system frequency suffers only small changes in order to employ rapid steady–state (i.e. constant frequency) methods to solve the network equations. When a power system is stable, this is a valid assumption. However, the loss of synchronism violates this assumption, meaning that our quantitative models break down at the stability limit.

Nevertheless, at the stability limit, the behaviour of the fast Fourier transform of the network transient response changes fundamentally and this qualitative observation suffices to discriminate stable from unstable cases.

6.3.2.4 A CRITERION FOR POWER SYSTEM TRANSIENT STABILITY

Based on the study of the Hydro–Québec system, the following criterion is proposed for evaluating the stability of transient stability simulations:

"If the polar plot of the Fourier transform of the system's transient response, as defined by equation (5–27), exhibits a sustained counterclockwise (CCW) phase angle rotation at any single network monitoring location, then this response corresponds to an unstable case."

To apply this criterion, in view of the limitations of frequency domain analysis in the case of extremely short unstable transient stability simulations (see 6.3.2.2), bus voltage must be monitored at a number of network locations. Of course, this constraint would no longer be necessary if unstable simulations were lengthened either by relaxing or disabling the angle criterion responsible for premature interruption of the transient stability simulation.

A final remark. As sustained CCW polar plot rotation is characteristic only of instability, this criterion is valid for all operating points, not only in the region near the stability limit as had been earlier postulated.

6.3.2.5 ADDITIONAL CONCERNS

The knowledgeable reader will have noticed that the polar plots published in the paper entitled "Frequency–Domain Behaviour of a Network Near Its Transient Stability Limit" [Marceau, Galiana, Mailhot & McGillis 1993] are not identical to those of the present chapter. In particular, the time–domain voltage waveforms of the paper's Figs. 5 and 6 are identical to Figs. 5–1 (a) and (b) of the present thesis, except for the voltage scale. This, in turn, implies that Figs.

7 and 8 of the paper should respectively be identical to those of Figs. 6–7 and 6–8 of the present thesis. However, the actual appearance of the polar plots is at times quite different from those presented in this thesis, particularly in relation to unstable cases, though the angle criterion proposed here is respected by all polar plots.

The paper's polar plots were generated using the MATRIX$_x$ software whereas all those presented in this thesis were obtained using MATLAB. The reason for this is that the Hydro–Québec Operations Planning department abandoned the use of MATRIX$_x$ in favour of MATLAB in late 1992, obliging me to continue with MATLAB after the paper had been written. The main feature of the paper's polar plots, obtained using MATRIX$_x$, is that unstable cases exhibited a sustained CCW rotation, including a spiralling–in characteristic with increasing frequency. The MATLAB *FFT*, though preserving CCW rotation for unstable cases, falls slowly into the origin. A possible explanation for these discrepancies is that the fast Fourier transform function supported by MATRIX$_x$ provides a "window" parameter which tends to "smooth out" the *FFT* at the expense of decreased resolution whereas the MATLAB *FFT* provides no such windowing parameter. It may be that, at the time the plots for the paper were made (i.e. in October 1992), the default settings in the Hydro–Québec MATRIX$_x$ environment applied such a window – without my knowledge – even though the manual explicitly states that this should not have been the case [Integrated Systems 1990]. Unfortunately, it is impossible to verify this hypothesis, or to explain it otherwise as MATRIX$_x$ is no longer available to me at this time.

The MATLAB fast Fourier transform function has been validated by comparison with those of known waveforms (i.e. square pulses, see Fig. 5–3), hence I am confident of the results given in the present chapter. Regardless of the particular software used to generate polar plots, polar plot rotational direction is the same.

6.4 Summary

Building on the signal energy approach presented in the previous chapter, the present chapter explores a network's frequency–domain behaviour near the transient stability limit.

In this chapter, a frequency–domain stability criterion is proposed, based on the observation that a network's behaviour near the stability limit is governed by the action of its dominant poles, as demonstrated in the previous chapter. Essentially, the criterion requires monitoring the Fourier transform of the network's transient response $R(\omega, P)$ at a number locations in the network. When generation P is varied and the network crosses the stability limit, the angle of $R(\omega, P)$ fundamentally changes its behaviour, passing from a clockwise (CW) to a counterclockwise (CCW) rotational behaviour, as monitored in at least one location. This is confirmed by results obtained from performing many stability limit searches on the Hydro–Québec system using ELISA. Once again, we see that these results hold for normal contingencies, notwithstanding the presence of many active, nonlinear voltage support elements in the network.

Such a criterion is useful in that the transient stability transfer limit determination process can now be mechanized without having to resort to heuristic or traditional acceptability criteria.

CHAPTER 7

CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

7.1 INTRODUCTION

The main focus of this thesis has been to mechanize the processes of dynamic security analysis in operations planning by means of software *frameworks*. In particular, the ELISA framework prototype was developed and subsequently tested in a typical utility production environment. This experience has showed that even a modest effort can result in large benefits for a utility, and there are many reasons for this. First, frameworks accelerate the execution of lengthy, time–intensive processes, enabling planners to be more strategic in the realization of their tasks. Second, frameworks permit a large increase in the number of degraded networks which can be studied explicitly, thereby increasing the number of transmission limits obtained by simulation. Third, the possibility to study far more degraded topologies reduces the uncertainty and guesswork often found in constructing security control strategies. Amazingly enough, this leads to two contradictory results: Less conservative security limits and greater system security.

With the advent of more powerful computer hardware and more effective simulation software, the industry trend towards on–line security strategies will inevitably cause frameworks, developed in and for the off–line environment, to gravitate towards the on–line environment. A corollary to this is that time–honored and effective off–line methodologies *should* be mechanized and validated in the off–line environment *before* they find their way into the on–line environment, so that present levels of security can be maintained. The extensive testing and subsequent utilization of the ELISA prototype in a typical operations planning environment is an important step in this direction.

Finally, as we have also seen in this thesis, the leverage afforded by such tools in the research environment can also contribute to augmenting our understanding of fundamental power systems phenomena, and significantly increase the effectiveness of existing tools and strategies.

7.2 SPECIFIC CONCLUSIONS ARISING FROM THIS WORK

From the perspective of the research performed in the course of this thesis, the following conclusions follow directly:

1. A taxonomy of the *processes* of dynamic security analysis in operations planning can be constructed using semantic nets, classes, objects, properties and rules, structured essentially around the goals of dynamic security analysis.

2. Such a taxonomy also defines the *language* of operations planners, a language of semantic *and* software generality in that the processes that it describes are independent of any particular network or commercial simulation software. As a direct consequence of this, the language can be viewed as a roadmap for constructing a very high–level *shell* environment or *framework:* unique bridges need only be built for each application program supported in a software *library*.

3. The ELISA framework prototype, built on the basis of this taxonomy, is found to have considerable flexibility and performance. In the utility environment, the prototype permits the execution of large–scale studies in a fraction of the time formerly required, and enables operations and system planners to work at a more strategic level. In the research environment, the ELISA prototype is also found to be extremely useful in the study and understanding of power systems phenomena, permitting unheard–of numbers of simulations to be performed in a relatively short period of time.

4 The signal energy E of the network voltage transient response, monitored at some user–defined location, is found to act as a barometer defining the relative severity of a contingency with respect to some parameter P, for instance power generation or power transfer. For a given contingency, as P is varied and the network approaches instability, signal energy increases smoothly and predictably towards an asymptote which defines the network's stability limit: This limit, in turn, permits us to compare the severity of different contingencies.

5. The effect of increasing P on the damping of the system can be explained using a Fourier transform approach. A simple function, derived from linear

systems theory, estimates network stability limits with surprising accuracy from only two or three sample simulations. In particular, using a formula based on three simulations (equation 5−23), less than 1% error results provided that 1) the monitoring location is close to the contingency location, 2) at least one of the simulations is within 3.5% of the stability limit and 3) the farthest simulation is within 16%. These results hold for all normal contingencies (i.e. faults with no change in topology, faults accompanied by subsequent loss of line or spontaneous loss of line), notwithstanding the presence of many active, nonlinear voltage−support elements (i.e. generators, synchronous condensers, SVCs, static excitation systems, etc.). Finally, signal energy behaviour permits us to understand and quantify the effectiveness of such long−used concepts as stability margin.

6. As implied in the previous point, signal energy monitoring location can impact the precision of the signal−energy−based stability limit estimates. Monitoring locations close to the contingency location give better stability limit estimates than distant monitoring locations.

7. A per unit signal energy function (i.e. a dimensionless signal energy) can also be defined whose slope establishes proximity to the stability limit in a general way and which more formally circumscribes the region of validity of the signal energy model. In particular, stability limit estimates obtained from the proposed signal energy model (equation 5−23) are found to be highly accurate (i.e. within 1% error) provided that the change of the per−unit signal energy E_{ij}' with respect to per−unit power P_i' is greater than or equal to 497, that is:

$$\frac{\Delta E_{ij}'(P_i')}{\Delta P_i'} \geq 497 \qquad (5-41)$$

Taking this a step further, one can assume that this slope also defines the entrance to the knee of the signal energy curve. Based on this assumption, one finds that the signal energy curve begins its rapid ascent at approximately 95.5% of the stability limit.

8. Building on these results, a strategy is proposed permitting rapid estimates of transient stability transfer limits from at most three stable transient sta-

bility simulations. This may eventually prove to be attractive for on–line dynamic security analysis strategies employing time–domain simulation tools.

9. Finally, a frequency–domain stability criterion is proposed for transient stability analysis, based on the observation that a network's behaviour near the stability limit is governed by the action of its dominant poles. Essentially, the criterion requires monitoring the Fourier transform of the network's transient response $R(\omega, P)$ at a number of locations in the network. When generation P is varied and the network crosses the stability limit, the angle of $R(\omega, P)$ fundamentally changes its behaviour, passing from a clockwise (CW) to a counterclockwise (CCW) rotational behaviour, in at least one location. Once again, these results hold for normal contingencies, notwithstanding the presence of many active, nonlinear voltage support elements in the network.

7.3 ADDITIONAL PHILOSOPHICAL CONSIDERATIONS ON OPERATING SYSTEMS, EXPERT SYSTEMS AND ARTIFICIAL INTELLIGENCE

Before closing this thesis, it is useful to dwell on a certain number of issues which naturally arise from the mechanization of human processes. Beyond the semantics of such words as "frameworks", "expert systems", or "artificial intelligence", there lies the reality of attempting to mimic humans in various areas of endeavour. In the context of the present research, while struggling with the difficulties of teaching computers to imitate humans, a certain number of observations were made which may be of value for the emergence of machine intelligence.

Let us recall that any software that can be executed under a particular operating system can be seen as *de–facto* integrated within this operating system. Programmable operating systems can be seen as going a step beyond, providing the means for creating ever higher–level software from elementary building blocks which are neither elementary nor meant to be used as "building blocks". Though expert–system shell technologies and artificial intelligence (AI) languages such as LISP and PROLOG attempt to provide environments for creating high–level strategies and interfacing with a variety of software including databases, such tools are inherently limited in that they are themselves resident

within and dependent on operating systems, and are not *in total control* of their environment. The operating system stands alone at the top of the software control pyramid.

In the long term, the quest for machine intelligence must strive towards total control of the software environment as intelligence is the highest form of software and such software can only reside at the top of the software pyramid. Consequently, operating systems technology must inevitably converge with knowledge representation and processing technologies. In the short term, a natural evolution of operating systems would be to offer object–oriented modelling and programming features, including the capacity to build rule bases, knowledge bases, databases, and to employ inference engines.

The strategy of controling one's environment invites additional comment. We have seen in Chapter 1 that a fundamental tenet of security (and survival) is to exercise the greatest possible control over one's environment. Let us now reason backwards: If the known strategy is to exercise control over one's environment, one infers that survival is a key goal. Transposing this to the sofware environment, if control over the environment constitutes a necessary condition for the emergence of intelligence, survival is therefore an inherent trait of such intelligence. One therefore reasons that the quest for survival is a *prerequisite* to intelligence, or alternatively, that intelligence arises as a *by–product* of some demonstrated capability to survive.

All of this might seem to diverge from the main topic were it not for the fact that a primary concern of this thesis is the dynamic *security* of power systems. We have already defined power system security as the art and science of power system *survival*. Pursuing our previous line of reasoning, one can view the power system security problem as an ideal problem in artificial intelligence. A software capable of assimilating its own survival to that of the power system, and incorporating all of the necessary attributes to ensure its survival (i.e. real–time monitoring, high–level reasoning, learning capability, access to simulation tools, ability to analyse results, ability to exercise control on its own computer environment and on the network, etc.) may well constitute an important research direction towards achieving a software–based intelligence ultimately greater than the sum of its parts.

7.4 RECOMMENDATIONS FOR FUTURE RESEARCH

7.4.1 ON FRAMEWORKS IN OPERATIONS AND SYSTEM PLANNING

A number of research directions can be identified regarding frameworks in operations and system planning, in particular to continue the work that has begun in the course of this thesis:

1. Viewing the ELISA object–oriented language as a programmable language, it may be interesting to address the issue of building a robust compiler of ELISA language scripts, capable of generating the appropriate ELISA code "made to measure" for a given script. The implementation of conditional branching structures, filters and the capability to execute large–scale sensitivity studies, as described in Chapter 3, are also very challenging projects, particularly from the perspective of compiler construction. Finally, it would be interesting to be capable of viewing the class–object network resulting from the execution of an ELISA script onscreen, as illustrated in Fig. 3–10, as the search process progresses.

2. A corollary to the preceding is to attempt to build a voice interpreter of ELISA commands, including the capability to furnish a corresponding verbal answer to the request. Research into this area may have interesting implications from the point of view of future EMS technology.

3. It would be useful to implement the transient stability transfer limit strategy described in Chapter 5 which proposes to find a stability limit from at most three stable simulations.

4. Though not explicitly mentioned in this thesis, occasions exist where the stability limit is approached by reducing rather than incrementing power generation or transfer in some part of the network. This implies that the sense of a search strategy can also be considered a variable. Integrating this feature to the ELISA prototype is also an interesting project.

5. As described in Chapter 3, it is possible to define heuristic dynamic security functions to answer dynamic security analysis queries based on the knowledge and data contained in an appropriate database. Such a heuristic shell

would be a fascinating project in that it would lay the foundations for generating the security limit tables ultimately used by the system operator, or for ensuring on-line security analysis from a database of cases.

6. The area of data validation, though crucial, is rarely addressed. It would be of significant interest to the industry to develop some means, heuristic or otherwise, of verifying the correctness of load flow and stability input data.

7. It may be of interest to develop an ELISA-like framework capable of driving large-scale processes employing EMTP-like software (i.e. Electro-Magnetic Transients Program) rather than transient stability software. Two approaches can be proposed: 1) To use EMTP instead of a transient stability software in ELISA or 2) to define a new taxonomy of processes related to the analysis of voltage-related phenomena.

8. It may also be of interest to develop an ELISA-like framework specifically designed for the system planning environment. This problem is characterized by a desired transmission capacity subject to predetermined system criteria: Hence, power system topology is the variable. A taxonomy of such processes would have to be defined before implementation issues could be addressed.

9. As suggested in the previous section, a software capable of assimilating its own survival to that of the power system, and incorporating all of the necessary attributes to ensure its survival (i.e. real-time monitoring, high-level reasoning, learning capability, access to simulation tools, ability to analyse results, ability to learn, ability to execise control on its own computer environment and on the network, etc.) might prove to be an extremely rich research direction from the perspective of future generations of EMS technology.

7.4.2 On Signal Energy, System Stability and Transfer Limits

Here again, a number of research directions can be identified to continue the work that has begun:

1. It would be of interest to establish whether or not signal energy is useful in establishing long–term stability limits. If not, an appropriate criterion should be researched in order to reduce the number of simulations required to find such limits.

2. It was shown in Chapter 5 that acceptable cases from an operations perspective are usually found in the knee region of the signal energy curve. It would be of interest to establish this correlation more formally and determine a criterion, either in terms of p.u. signal energy slope or power in p.u. of the stability limit, which will define case acceptability more generally.

3. It would be of interest to explore signal energy from the perspective of different network quantities (i.e. frequency, power flow, reactive power flows, etc.) for transient stability transfer limit determination.

4. It would be of interest to further explore signal energy from the perspective of monitoring location. Would it be significant to sum the signal energies of the voltage at all the system buses, at all of the buses of only one voltage level, or at a small number of buses chosen on the basis of some criterion as yet to be determined? Would better stability limit estimates result?

5. It would be of interest to further study the impact of voltage support technologies (i.e. SVCs, SCs, static excitations, etc.) on signal energy curvature and determine if any new information can be extracted from this approach, in particular for system planning purposes.

6. It may be possible to improve the signal energy formula of equation (5–27) by considering the changes in dominant–pole frequency with respect to increasing power.

7. It would be of interest to find some correlation between the steady–state transfer limit of a power system topology, the transient stability transfer limit of the same system subjected to a fault with no change in topology, the transient stability transfer limit of the system degraded by one line and subjected to the same fault, and the transient stability transfer limit of the undegraded system subjected to a fault and loss of line.

8. It would be useful to confirm the validity of the proposed stability criterion by verifying it on simulations of other power systems.

9. It may be worthwhile to explore the applicability of wavelet theory to the stability evaluation of transient stability simulations, particularly in relation to short simulation times.

10. Is it possible to define a second equation including the signal energy constants enabling stability limits to be defined from a single transient stability simulation?

11. Is it possible to extend signal energy analysis and stability limit estimation such that on–line system data might be used to determine stability limits without simulation?

Aho, A., Kernighan, B.W., Weinberger, P.J., The AWK Programming Language, Addison–Wesley, Reading, Mass., 1988.

Akimoto, Y., Tanaka, H., Yoshizawa, J., Klapper, D.B., Price, W.W., Wirgau, K.A., "Application of Expert Systems to Transient Stability Studies", Second Symposium on Expert Systems Applications to Power Systems, Seattle, July, 1989.

Anderson, G., Anderson, P., The UNIX C–Shell Field Guide, Prentice–Hall, Englewood Cliffs, N.J., 1986.

Anderson, P.M. and Fouad, A.A., Power System Control and Stability, Iowa State University Press, Ames, Iowa, 1977.

Athay, T., Podmore, R., Virmani, S., "A Practical Method for Direct Analysis of Transient Stability", IEEE Tr. on Power Apparatus and Systems, PAS–98, No. 2, 1979, pp. 573–584.

Avramovic, B., Fink, L.H., Artificial Intelligence in Power System Voltage Control (Phase 1), Research report performed under National Science Foundation award number ISI–9060250 (U.S.A.), September 30, 1991.

Balu, N., Bertram, T., Bose, A., Brandwajn, V., Cauley, G., Curtice, D., Fouad, A., Fink, L., Lauby, M.G., Wollenberg, B.F., Wrubel, J.N., "On–Line Power System Security Analysis", Proceedings of the IEEE, Vol 80, No. 2, February 1992, pp. 280.

Brandwajn, V., "Efficient Bounding Method for Linear Contingency Analysis", IEEE Tr. on Power Systems, Vol. 3, No. 1, Feb. 1988, pp.38–43.

Brandwajn, V., Lauby, M., "Complete Bounding Method for AC Contingency Screening", IEEE Tr. on Power Systems, Vol. 4, No. 2, May 1989, pp. 724–727.

Byerly, R.T. and Kimbark, E.W. (editors), Stability of Large Electric Power Systems, IEEE Press, Piscataway, N.J., 1974.

Carr, C., "The Man of Silence", Experience of War, (Editor Cowley, R.), W.W. Norton & Co., New York, 1992, p. 235.

Carpentier, J., "Towards a Secure and Optimal Automatic Operation of Power Systems", (Keynote Address) Proceedings of the 1987 PICA Conference, pp. 2–37.

Carpentier, J.L., Gillon, A., Jegouzo, Y., Candre, A., Caraman, F., Pellen, C., Tournebise, P., "Multimethod Optimal Power Flows at Électricité de France", Proceedings of the IFAC Symposium on Power System and Power Plant Control, August 22–25, 1989, pp. 129–134.

Chandrasekaran, B., "Expert Interview: Allen Newell", IEEE Expert, Vol. 8, No. 3, June 1993, pp. 5–12.

Chai, J. S., Bose, A., "Bottlenecks in Parallel Algorithms for Power System Stability Analysis", 1992 IEEE Winter Power Meeting, New York, N.Y., 1992, 92 WM 285–7 PWRS.

Chiang, H.D., Tong, J., Miu, K.N., "Predicting Unstable Modes in Power Systems: Theory and Computations", 1992 IEEE Summer Power Meeting, 92 SM 591–8.

Churchill, R.V., Brown, J.W., Verhey, R.F., Complex Variables and Applications, McGraw–Hill Inc., New York, 1974.

Cicoria, R., Migliardi, P., Marannino, P., "Knowledge Based Methodologies versus Conventional Optimisation Algorithms: Applications to Reactive Power Compensation", Proceedings of the 11th Power Systems Computation Conference, Aug. 30–Sept. 3, 1993.

Cihlar, T.C., Wear, J.H., Ewart, D.N., Kirchmayer, L.K., "Electric Utility System Security", Proceedings of the American Power Conference, Vol. 31, pp.97–109, 1969.

Cleary, T. (translator and editor), Zhuge Liang & Liu Ji (authors) Mastering the Art of War: Zhuge Liang's and Liu Ji's Commentaries on the Classic by Sun Tzu, Shambhala Publications Inc., Boston, 1989, pp. 1–2, 45.

Corwin, J.L., Miles, W.T., Impact Assessment of the 1977 New York City Blackout, Systems Engineering for Power, Div. of Electrical Energy Systems, U.S. Dept. of Energy, July, 1978.

Crevier, D., Schweppe, F.C. "The Use of Laplace Transforms in the Simulation of Power System Frequency Transients", IEEE Transactions on Power Apparatus and Systems, vol. PAS-94, no. 2, March/April 1975, pp. 236–241.

Crevier, D., "A Simulation–Based Separation Surface Method for Power System Stability Assessment", IEEE PES Winter Meeting, New York, NY, Jan/Feb 1978, paper A 78 226–3.

Crainic, E.D., Horisberger, H.P., Do, X.D., Mukhedkar, D., "Power Network Observability: The Assessment of the Measurement System Strength", IEEE Tr. on Power Systems, Vol. 5, No. 4, pp. 1267–1285.

Crow, M.L., Tylavsky, D.J., Bose, A., "Concurrent Processing in Power System Analysis", Analysis and Control System Techniques for Electric Power systems, Vol. 42, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 1–56.

Dénommé, F., Vincelette, M., "Security Monitoring of the Hydro–Québec Power System", Proceedings of the Pan–American Congress of Mechanical, Electrical and Allied Engineering Branches, Buenos Aires, Argentina, Oct. 1–5, 1984.

Desoer, C.A., Notes for a Second Course on Linear Systems, D. Van Nostrand Company, New York, 1970.

Dillon, T.S., Laughton, M.A., (Editors), Expert System Applications in Power Systems, Prentice–Hall International, London, U.K., 1990.

Dy Liacco, T.E., "The Adaptive Reliability Control System", IEEE Tr. on Power

App. and Systems, Vol. PAS-86, pp. 517-531, May 1967.

Dy Liacco, T.E., Control of Power Systems via the Multi-Level Concept, Case Western Reserve University, Systems Research Center, Report SRC-68-19, June 1968.

Dy Liacco, T.E., "Real-Time Computer Control of Power Systems", Proceedings of the IEEE, Vol. 62, pp. 884-891, July 1974.

Dy Liacco, T.E., "System Security: The Computer's Role", IEEE Spectrum, Vol. 15, pp. 43-50, June 1978.

Ejebe, G.C., Wollenberg, B.F., "Automatic Contingency Selection", IEEE Tr. on Power Apparatus and Systems, Vol. PAS-98, No. 1, Jan./Feb. 1979, pp. 97-109.

Elgerd, O.I., Control Systems Theory, McGraw-Hill Book Company, New York, N.Y., 1967.

Elgerd, O.I., Electric Energy Systems Theory: An Introduction, McGraw-Hill book Company, New York, N.Y., 1971.

El-Sharkawi, M.A., Marks, R.J. II, Proceedings of the First International Forum on Applications of Neural Networks to Power Systems, Seattle, Washington, 1991.

Ewart, D.N., "Why and Wherefores of Power System Blackouts", IEEE Spectrum, Vol 15, 1978.

Fink, L.J., Carlsen, K., "Operating Under Stress and Strain", IEEE Spectrum, Vol. 15, March 1978.

Fishl, R., Mercede, F., Wu, F.F., Chiang, H.-D., "A Comparison of Dynamic Security Indices Based on Direct Methods", Electrical Power and Energy Systems, Vol. 10, No. 4, October 1988, pp. 210-231.

Fouad, A.A. (Chairman), "Dynamic Security Assessment Practices in North

America", Report by IEEE Working Group on Dynamic Security Assessment, Power Systems Engineering Committee, IEEE Tr. on Power Systems, Vol. 3, No. 3, Aug. 1988, pp. 1310–1321.

Fouad, A.A., Vittal, V., "The Transient Energy Function Method", Electrical Power and Energy Systems, Vol. 10, No. 4, October 1988, pp. 233–246.

Fouad, A.A., Vittal, V., Power System Transient Stability Analysis Using the Transient Energy Function Method, Prentice–Hall, 1991.

Fouad, A.A., Venkataraman, S., Davis, J.A., "An Expert System for Security Trend Analysis of a Stability–Limited Power System", IEEE Tr. on Power Systems, Vol. 6, No. 3, Aug. 1991, pp. 1077–1084.

Frederick, D.K. and Carlson, A.B., Linear Systems in Communications and Control, John Wilesy and Sons, New York, 1971, pp. 218–261, 100–109, 58–64.

Gaba, G., Audette, L., Guillemette, F., Lafrance, F., RP600 logiciel de répartition de puissance: manuel d'utilisation, Hydro–Québec, Vice–présidence technologie et IREQ, Avril 1993.

Galiana, F.D., "Bound Estimates of the Severity of Line Outages in Power System Contingency Analysis and Ranking", IEEE Tr. on Power Apparatus and Systems, Vol. PAS–103, No. 9, Sept. 1984., pp. 2612–2624.

Galiana, F.D., Schweppe, F.C., Glavitsch, H., "Synthesis of Security Monitoring Schemes in Power Systems", Proceedings of the 4th IFAC/IFIP International Symposium on Digital Computer Applications to Process Control, Zurich, Switzerland, Part II, March 19–22, 1974.

Galiana, F.D., Zeng, Z.C., "Analysis of the Load Flow Behaviour Near a Jacobina Singularity", IEEE Tr. on Power Systems, Vol. 7, No. 3, August 1992, pp. 1362–1369.

Galiana, F.D., McGillis, D., Marin, M., "Expert Systems in Transmission Plan-

ning", Proceedings of the IEEE, Vol. 80, No. 5, May 1992.

Gao, B., Morison, G.K., Kundur, P., "Voltage Stability Evaluation Using Modal Analysis", IEEE Tr. on Power Systems, Vol. 7, No. 4, Nov. 1992, pp. 1529–1542.

Genesereth, M.R., Nilsson, N.J., Logical Foundations of Artificial Intelligence, Morgan Kaufmann Publishers, Inc., Palo Alto, CA, 1988, pp. 36–41.

Halpin, T.F., Fischl, R., Fink, P., "Analysis of Automatic Contingency Selection Algorithms", IEEE Tr. on Power Apparatus and Systems, Vol. PAS–103, No. 5, May 1984, pp. 938–945.

Hingorani, N.G., "Flexible AC transmission", IEEE Spectrum, Vol. 30, No. 4, April 1993, pp. 40–45.

Hiskens, I.A., Hill, D.J., "Incorporation of SVCs into Energy Function Methods", 1991 IEEE Summer Power Meeting, 91 SM 424–2.

Hydro–Québec, Design Criteria and Practices of Hydro–Québec's Power Transmission System, Dec. 1990.

Horowitz, S.H. (Editor), Protective Relaying for Power Systems, IEEE Press, New York, N.Y., 1980.

IEEE Task Force on Terms & Definitions, "Proposed Terms & Definitions for Power System Stability", (System Dynamic Performance Subcommittee, Power System Engineering Committee), IEEE Tr. on Power Apparatus and Systems, Vol. PAS–101, No. 7, July 1982.

IEEE PSRC Working Group K12 Report, Voltage Stability of Power Systems: Concepts, Analytical Tools, and Industry Experience, IEEE Publishing Services, Piscataway, N.J., 1991.

Integrated Systems Inc., MATRIX$_x$ Core, Integrated Systems Inc., Santa Clara, California, 1990.

Irissari, G.D., Levner, D., "Automatic Contingency Selection for On-Line Security Analysis - Real-Time Tests", IEEE Tr. on Power Apparatus and Systems, Vol. PAS-98, No. 5, Sept./Oct. 1979, pp. 1552-1559.

Kaplan, G., "How Users View Frameworks", IEEE Spectrum, Vol. 29, No. 11, November 1992.

Kimbark, E.W., Direct Current Transmission (Volume 1), Wiley-Interscience, New York, 1971, Chapter 5.

Langevin, M., La stabilité des réseaux hydro-électriques par les règlages: présentation d'une méthode de régulation modale adaptative, thèse de Ph. D. présentée à l'École Centrale de Lyon, décembre 1987.

Lauby, M.G., Mikolinnas, T.A., Reppen, N.D., "Contingency Selection of Branch Outages Causing Voltage Problems", IEEE Tr. on Power Apparatus and Systems, Vol. PAS-102, No. 12, December 1983, pp. 3899-3904.

Lauby, M.G., "Evaluation of a Local DC Flow Screening Method for Branch Contingency Selection of Overloads", IEEE Tr. on Power Systems, Vol. 3, No. 3, August 1988, pp. 923-928.

Leite da Silva, A.M., Endrenyi, J., Wang, L., "Integrated Treatment of Adequacy and Security in Bulk Power System Reliability Evaluations", IEEE Tr. on Power Systems, Feb. 1993, p. 275.

Lim, E.-P., Cherkassky, V., "Semantic Networks and Associative Databases: Two Approaches to Knowledge Representation and Reasoning", IEEE Expert, Vol. 7, No. 4, August 1992, pp. 31-40.

Little, J.N., Shure, L., Signal Processing Toolbox For Use With MATLAB: User's Guide, The Mathworks, Inc., Natick, Mass., 1992.

Liu, C.-C., Dillon, T.S., "State of the Art", Expert System Applications in Power Systems, (Dillon, T.S., Laughton, M.A., editors), Prentice-Hall, London (U.K.), 1990.

Liu, C.–C., "Knowledge–Based Systems in Power Systems: Applications and Development Methods", Tr. IEE Japan, Vol. 110–B, No. 4, 1990.

Loparo, K.A., Abdel–Malek, F., "A Probabilistic Approach to Power System Security", IEEE Tr. on Circuits and Systems, Vol. 37, No. 6, June 1990, pp. 787–798.

Malik, O.P., Zhang, Y., "Application of Neural Nets to the Excitation Control of a Generating Unit", presented at the Engineering And Operating Division Meeting, Canadian Electrical Association, Vancouver, March 1992.

Mansour, Y., Kundur, P., "Voltage Collapse: Industry Practices", Analysis and Control System Techniques for Electric Power systems, Vol. 42, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 111–162.

Mailhot, R., Gauthier, J., Vanier, G., Étude de la stabilité à long terme de la tension du réseau d'Hydro–Québec, Hydro–Québec, Direction Exploitation, Service Réseau, 1993.

Marceau, R.J., Mailhot, R., Galiana, F.D., "A Generalized Shell for Dynamic Security Analysis in Operations Planning", 1992 IEEE Summer Power Meeting, 92 SM 435–8.

Marceau, R.J., Galiana, F.D., Mailhot, R., Dénommé, F., McGillis, D., "Fourier Methods for Estimating Power System Stability Limits", 1993 IEEE PICA Conference, Phoenix, Arizona, May, 1993.

Marceau, R.J., Galiana, F.D., Mailhot, R., McGillis, D., "Frequency–Domain Behaviour of a Network Near Its Transient Stability Limit", 1993 11th Power Systems Computation Conference, Avignon, France, Aug./Sept., 1993.

McGillis, D., Galiana, F.D., Loud, L., Marceau, R.J., "The Influence of the Choice of Criteria on System Design", Proceedings of the 9th Conference on the Electric Power Supply Industry (CEPSI), Hong Kong, Nov. 23–27, 1992.

McGillis, D., Galiana, F.D., Manoliu, R., "An Integrated Package of Expert Systems For Transmission Planning", Proceedings of the International Power Engineering Conference (IPEC), Singapore, March 18–19, 1993.

Mikolinnas, T.A., Wollenberg, B.F., "An Advanced Contingency Selection Algorithm", IEEE Tr. on Power Apparatus and Systems, Vol. PAS–100, No. 2, February 1981, pp.608–617.

Minsky, M., "A Framework for Representing Knowledge", in P. Winston (Ed.) The Psychology of Computer Vision, McGraw–Hill, 1975.

National Energy Board of Canada, Technical Features of Security of Electric Power System Supply, (report for the information of the electric power utilities of Canada; Stead, R.A., Secretary), May, 1967.

Neuron Data, Nexpert Object: Functional Description, Neuron Data documentation, Palo Alto, 1991.

North American Reliability Council (NERC), Reliability Concepts in Bulk Power Electric Systems, 1985.

Nourmoussi, M.A., Crevier, D., "Multi–Machine Power System Stabilizer Optimization by Maximization of the Transient Stability Domain", IEEE PES Summer Meeting, Vancouver, BC, July 1979, paper A 79 467–2.

Oh, S.–Y., "A Pattern Recognition and Associative Memory Approach to Power System Security Assessment", IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC–16, No.1, Jan./Feb. 1986, pp. 62–72.

Pai, M.A., Energy Function Analysis for Power System Stability, Kluwer Academic Publishers, 1989.

Pang, C.K., Prabhakara, F.S., El–Abiad, A.H., Koivo, A.J., "Security Evaluation in Power Systems Using Pattern Recognition", IEEE PES Summer Meeting, Vancouver, July 1973, pp. 969–976.

Papoulis, A., The Fourier Integral and its Applications, McGraw–Hill, New York, 1962.

Parsaye, K., Chignell, M., Expert Systems for Experts, John Wiley & Sons, Inc., New York, 1988, p. 141.

Pasquier, S. "Il a rêvé, il a fait Sony...", Japon: le triomphe et les complexes, (collection Les cahiers de l'Express), Groupe Express S.A., Paris, janvier 1993 p. 74.

Phadke, A.G., Thorp, J.S., "Improved Control and Protection for Power Systems Trough Synchronized Phasor Measurements", Analysis and Control System Techniques for Electric Power systems, Vol. 43, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 335–375.

Ribbens–Pavella, M., Murphy, P.G., Horward, J.L., Carpentier, J., "On–Line Transient Stability Assessment and Contingency Analysis", CIGRE, 1982, Report No. 32–19.

Saitoh, H., Toyoda, J., Kobayashi, Y., "A New Index Extracted From Line Flow Fluctuations to Evaluate Power System Damping", 1991 IEEE Winter Power Meeting, 91 WM 208–9.

Saucedo, R., and Schiring, E.E., An Introduction to Continuous and Digital Control Systems, Collier–Macmillan, London, 1968, Chapters 7 and 10.

Sauer, P.W., Pai, M.A., "Modeling and Simulation of Multimachine Power System Dynamics", Analysis and Control System Techniques for Electric Power systems, Vol. 43, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 1–60.

Schneiderman, B., Designing the User Interface, Addison–Wesley, Reading, Mass., 1986, Chapter 6.

Shoults, R.R., Jativa, J.A., "Real–Time Power System Control: Issues related to Variable Nonlinear Tie–line Frequency Bias for Load Frequency Control",

Analysis and Control System Techniques for Electric Power systems, Vol. 43, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 377–406.

Stott, B., Alsac, O., Monticelli, A.J., "Security Analysis and Optimization", Proceedings of the IEEE, Vol 75, No. 12, Dec. 1987, pp. 1623–1644.

Sun–Tzu, The Art of Warfare, (translated by Roger T. Ames), Ballantine Books, New York, 1993, see pp. 82–83, 104 on "strategic advantage".

Talukdar, S., Cardozo, E., "Building Large–Scale Software Organizations", Expert Systems for Engineering Design (Michael D. Rychener editor), Academic Press Inc., Boston, 1988, pp. 241–256.

Taub, H. and Schilling, D.L., Principles of Communication Systems, McGraw–Hill Inc., New York, 1971.

Thorp, J.S., Phadke, A.G., "Computer Relaying in Power Systems", Analysis and Control System Techniques for Electric Power systems, Vol 44, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 1–58.

Tinguely, C., Système–expert pour l'analyse de sécurité d'un réseau de transport d'énergie électrique, (Ph.D. thesis, No. 1089), Ecole Polytechnique fédérale de Lausanne, 1992.

Torre, W.V., Bushner, R.E., "A Unique Method of Evaluation of System Damping in Simulations of Large Power Systems", IEEE Tr. on Power Systems, Vol. PWRS–2, No. 1, pp. 119–122, February 1987.

Trotignon, M., Counan, C., Maury, F., Lesigne, J.F., Bourgin, F., Tesseron, J.M., Boisseau, J., "Plan de défense du réseau THT français contre les incidents généralisés: dispositions actuelles et perspectives d'évolution", rapport CIGRE 39–306 présenté au COMITÉ 39 de CIGRE, session de Paris août–septembre 1992.

Valette, A., Lafrance, F., Lefebvre, S., Radakovitz, L., ST600 programme de stabilité: manuel d'utilisation version 701, Hydro–Québec, Vice–présidence

technologie et IREQ, janvier 1987.

Venkata, S.S., Damborg, M.J., Jampala, A.K., "Power System Protection: Software Issues", Analysis and Control System Techniques for Electric Power systems, Vol 42, (C.T. Leondes editor), Academic Press, San Diego, 1991, pp. 57–110.

Vittal, V., (Chairman), "Transient Stability Test Systems for Direct Stability Methods", IEEE Committe Report, 1991 Winter Power Meeting, 91 WM 224–6.

Wang, Y., Hill, D.J., Middleton, R., Gao, L., "Transient Stability Enhancement and Voltage Regulation of Power Systems", 1992 IEEE Winter Power Meeting, New York, N. Y., January 1992, 92 WM 113–1 PWRS.

Waterman, D.A., A Guide to Expert Systems, Addison–Wesley, Reading, Mass., 1986, pp. 70–77.

Webster, N., et al., Webster's Ninth New Collegiate Dictionary, Merriam–Webster Inc., Springfield, Mass., 1987.

Wehenkel, Van Cutsem, T., L., Pavella, M., "An Artificial Intelligence Framework for On–line Transient Stability Assessment of Power Systems", IEEE Tr. on Power Systems, Vol. 4, No. 2, May 1989.

Wehenkel, L., Pavella, M., Euxibie, E., Heilbronn, B., "Decision Tree Based Transient Stability Method: A Case Study", 1993 IEEE Winter Power Meeting, Columbus, Ohio, January 1993, 93 WM 235–2 PWRS.

Westinghouse Electric Corporation, Electrical Transmission and Distribution Reference Book, Fourth edition, Westinghouse Electric Corporation, 777 Penn Center Boulevard, Pittsburgh, Pennsylvania 15235, 1950.

Wood, A.J., Wollenberg, B.F., Power Generation, Operation and Control, John Wiley & Sons, New York, N.Y., 1984, Chapter 11.

Wu, F.F., Tsai, Y.-K., "Probabilistic Dynamic Security Assessment of Power Systems: Part 1 – Basic Model", IEEE Tr. on Circuits and Systems, Vol. CAS-30, No. 3, March 1983, pp. 148-159.

Xue, Y., Van Cutsem, T., Ribbens-Pavella, M., "A Simple Direct Method for Fast Transient Stability Assessment of Large Power Systems", IEEE Tr. on Power Systems, Vol 3, No. 2, May 1988, pp.400-412.

Xue, Y., "Preliminary Design of A Hybrid System for Transient Security Analyses of Power Systems", Symposium on Expert Systems Applications to Power Systems, Helsinki, August, 1988.

Xue, Y., Van Cutsem, T., Ribbens-Pavella, M., "Extended Equal Area Criterion: Justifications, Generalizations, Applications", IEEE Tr. on Power Systems, Vol 4, No. 1, Feb. 1989, pp. 44-51.

Xue, Y., "A Planning Decision Support Expert System for Transient and Dynamic Security of Power Systems", Second Symposium on Expert Systems Applications to Power Systems, Seattle, July 17-10, 1989, pp. 218-223.

Xue, Y., Wehenkel, L., et al., "Extended Equal Area Criterion Revisited", 1991 IEEE Summer Power Meeting, 91 SM 422-6.