

# Metrics and Laws of Software Evolution - The Nineties View

M M Lehman

J F Ramil

P D Wernick

Department of Computing

Imperial College of Science, Technology and  
Medicine

London SW7 2BZ

tel: +44 (0)171 594 8214

fax: +44 (0)171 594 8215

e-mail: {mml,jcf1,pdw1}@doc.ic.ac.uk

URL: <http://www-dse.doc.ic.ac.uk/~mml/feast1/>

D E Perry

Bell Laboratories, Murray Hill, NJ 07974

+1 908 582 2529

dep@research.bell-labs.com

W M Turcki

Institute of Informatics

Warsaw University

Warsaw 02-097

+48 22 658 3522

wmt@mimuw.edu.pl

## Abstract

The process of E-type software development and evolution has proven most difficult to improve, possibly due to the fact that the process is a multi-input, multi-output system involving feedback at many levels. This observation, first recorded in the early 70s during an extended study of OS/360 evolution, was recently captured in a FEAST hypothesis; a hypothesis being studied in on-going two-year project, FEAST/1. Preliminary conclusions based on a study of a financial transaction system, FW, are outlined and compared with those reached during the earlier OS/360 study. The new analysis supports, or better does not contradict, the laws of software evolution, suggesting that the 1970s approach to metric analysis of software evolution is still relevant today. It is hoped that FEAST/1 will provide a foundation for mastering the feedback aspects of the software evolution process, opening up new paths for process modelling and improvement.

Keywords: Software:-process, evolution, process metrics, dynamics and improvement; Lehman's laws

## 1 Introduction

A 1968 study of the IBM software programming process<sup>1</sup> [leh69,85] led, *inter alia*, to metric based studies of OS/360<sup>2</sup> [bel72,leh74,85] and other systems [leh80b,kit82]. Analysis of data relating ultimately to some 26 of OS/360 releases and sub-releases, identified and ordered by their release sequence number *rsn* [cox66], yielded insights into various aspects of its evolutionary

<sup>1</sup> Unless otherwise stated, all references in this paper to *process* refer to both *ab initio* development and to subsequent system maintenance, that is, to the process of all aspects of system evolution [leh85].

<sup>2</sup> References in this paper to the IBM OS/360 system, refer to both that system and its successor OS/370.

trends. An example of the study output is provided by figure 1.

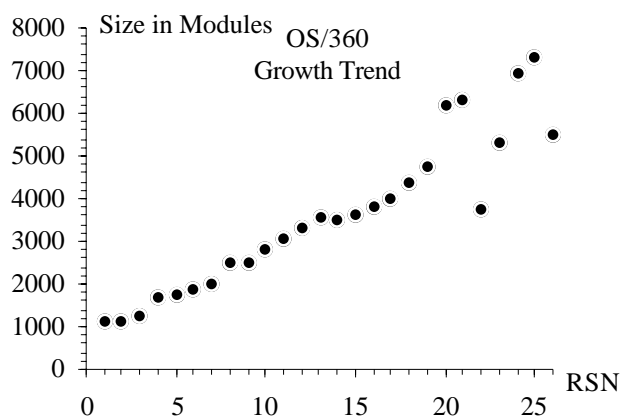


Fig. 1 OS/360 growth trend by rsn

Up to release  $rsn_{21}$ , the OS/360 growth trend illustrated by figure 1 may be interpreted as a small *ripple* superimposed on otherwise smooth growth. This pattern is reminiscent of traces generated by self-regulating and self-stabilising systems with both positive and negative feedback [bel72, leh78]. The behaviour thereafter may be interpreted as a sign of instability induced by excessive positive feedback; rapid functional evolution which led to a fission process; the transition from OS/360 to VS1 and VS2. Alternatively it may be interpreted as chaos-like behaviour. Either interpretation suggests that further prediction based on earlier behaviour is uncertain.

It was these observations that first suggested that software evolution processes are and must be treated as being feedback driven and constrained systems [leh74,78]. Subsequently, preliminary examination of data on several

non-IBM systems repeated many of the earlier observations [leh77, bel78].

The feedback theme was also applied by Abdel-Hamid and Madnick [abd91] in their work on the use of *system dynamics* in modelling software project management issues. This approach and related modelling techniques originated in the seminal work of Forrester and his colleagues at M.I.T [for61,70]. More recently it has been applied to the study of aspects of software process improvement, eg. [abd91, mcg93, wae94, mad96].

The 1970s study, exemplified by figure 1, revealed regularity unexpected for variables the indirect consequences of human decision but not directly or intentionally controlled. It could and was, therefore, captured in a series of statements abstracting the observed behaviour. Such abstraction involves a domain larger than the realm of software technology *as normally understood*. In fact, it includes software technology as a sub domain. From the point of view of software engineering, such statements must, therefore, be accepted as an external regulating and constraining force. To overcome them requires expertise in organisational dynamics, management, sociology etc., not just software technology. Thus they were subsequently referred to as *laws* of software evolution [leh74,78,80,85]. The laws, as proposed then [leh74] and subsequently amended [leh78,80], are summarised in table 1 with column 1 indicating the year when each was first published.

The laws could, however, only be applied with any degree of confidence to the domains to which the data related. Generalisation depended on obtaining confirming evidence from other systems and organisations. If achieved, such generalisation would provide a theoretical and practical base and framework for the evolution of *E-type* programs, that is, software solving a problem or addressing an application in the real world [leh80b].

## 2 Process improvement

In recent years many businesses, and the software industry in particular, have developed a strong interest in and commitment to disciplined process improvement. More and more business processes are, however, dependent on software generated information. They are driven and controlled by computers and software, probably including *E-type legacy systems* that may have been operational for many years. Now specification and design of such systems requires assumptions about the intended application and its operational domain. These, in turn, will be reflected in the software. Subsequently, installation and operation of the system together with exogenous change will invalidate some of the embedded assumptions [leh89]. The system must, therefore, be continually updated to maintain their validity and adapt to changed circumstances. Business and software process improvement are strongly linked and interdependent [leh97].

No.	Brief Name	Law
I 1974	Continuing Change	<i>E-type</i> systems must be continually adapted else they become progressively less satisfactory.
II 1974	Increasing Complexity	As an <i>E-type</i> system evolves its complexity increases unless work is done to maintain or reduce it.
III 1974	Self Regulation	<i>E-type</i> system evolution process is self regulating with distribution of product and process measures close to normal.
IV 1980	Conservation of Organisational Stability (invariant work rate)	The average effective global activity rate in an evolving <i>E-type</i> system is invariant over product lifetime.
V 1980	Conservation of Familiarity	As an <i>E-type</i> system evolves all associated with it, developers, sales personnel, users, for example, must maintain mastery of its content and behaviour [leh80a] to achieve satisfactory evolution. Excessive growth diminishes that mastery. Hence the average incremental growth remains invariant as the system evolves.
VI 1980	Continuing Growth	The functional content of <i>E-type</i> systems must be continually increased to maintain user satisfaction over their lifetime.
VII 1996	Declining Quality	The quality of <i>E-type</i> systems will appear to be declining unless they are rigorously maintained and adapted to operational environment changes.
VIII 1996	Feedback System (first stated 1974, formalised as law 1996)	<i>E-type</i> evolution processes constitute multi-level, multi-loop, multi-agent feedback systems and must be treated as such to achieve significant improvement over any reasonable base.

**Table 1 Laws of software evolution**

The present paper focuses on software process improvement. The fact is that the software industry has been seeking improvement of the software development and maintenance process for many years [wil51,leh96]. Academic and industrial effort has yielded incremental improvement, through the introduction of new languages, formalisation, improved methods, mechanised support (CASE), new programming paradigms and so on. Nevertheless, the industrial track record raises the question why, despite so many advances, the global software development process from conception to use is still so often marred; why satisfactory functionality, performance and quality is only achieved over a lengthy evolutionary process, why software maintenance never ceases until a system is scrapped, why software is still generally regarded as the weakest link in the development of computer-based systems [leh94,96].

Explanations for individual failures can always be found. This paper summarises a more general approach arising from recognising that development and evolution processes for *E*-type systems are intrinsically feedback systems [leh94]. The remainder of the paper reports preliminary results from an investigation of this hypothesis.

### 3 FEAST (*Feedback, Evolution And Software Technology*) and FEAST/1

Some years ago, the realisation that feedback<sup>3</sup> in the software process could explain the difficulties encountered in achieving its global improvement, led to the formulation of a FEAST hypothesis [fea94,leh94]:

*As complex feedback systems, E-type software processes evolve strong system dynamics and with it the global stability characteristics of other such systems. Consequent stabilisation effects are likely to constrain efforts at process improvement.*

More recently the hypothesis was restated in the following terms [leh96c]:

*As for other complex feedback systems, the dynamics of the real world software development and evolution processes will possess a degree of autonomy and global stability.*

Both versions of the hypothesis include a number of assertions [leh96a] but these are not further discussed here.

The hypothesis and its implications were examined over a period of time by a FEAST core group, a subgroup of the present authors. Its deliberations [leh96b] led to three workshops held at Imperial College during 1994/5 [fea94,95]. The objectives were to expose the ideas to a wider group of people interested in the software process, to seek the objective criticism of experts and, in general, to explore the hypothesis and its implications.

<sup>3</sup> The term feedback may be interpreted in several different ways. This theme has been discussed in [leh96b].

The EPSRC<sup>4</sup> proposal that resulted from these discussions was entitled FEAST/1 [leh96c], the "/1" in the title indicating that this two year, 3 person project was to be seen as a first step in a longer and more widespread investigation. The proposal was approved in March 1996 and the resultant project commenced formal investigations in October 1996 in collaboration with ICL plc, Logica plc, Matra-BAe Dynamics plc and two groups within the UK Ministry of Defence. The stated objectives [leh96c] were to:

- provide objective evidence that feedback phenomena and the consequent system dynamics have substantial impact in the software process
- demonstrate that the phenomena can be exploited in both managing and improving industrial processes
- produce justification for a wider and more substantial study based on the feedback perspective.

The two year investigation will seek to identify and characterise the feedback mechanisms active in the process, their impact on process characteristics and methods for applying the understanding gained to improve the respective processes. It expects to demonstrate (or otherwise) the feedback behaviour of some widely different systems. Three approaches are being employed:

- a *black box* approach is studying quantitative data from a number of industrial software processes to identify patterns in the evolution of the respective systems. The data will be analysed in a search for footprints of dynamical behaviour and feedback control
- a *white box* approach aims at the construction and enactment of system dynamics models of individual processes. These will reflect feedback mechanisms, their properties and their impact on the global process
- a third approach, not included in the original proposal and not formally part of FEAST/1, is exploring the use of multi-agent systems [mca95] to model the selected processes and to evaluate proposed improvements.

Full investigation of the hypothesis and, if upheld, of means for its exploitation, is not straightforward. As previously discussed [leh96b], difficulties arise from several factors. The processes being investigated are likely to include tens, if not hundreds, of forward paths and feedback loops. A simulation approach such as the system dynamics technique referred to in section 1 is, therefore more appropriate tool for the investigation than the analytical tools of control theory. Moreover, the processing and control mechanisms associated with these loops involve people, individually and in groups as managers or implementors. All observe, interpret, communicate, decide and act or refrain from acting on the basis of their overall perception, their instructions and, consciously or otherwise, their inclinations, experience and biases. Much of the feedback control is unplanned or even unconscious. Some, at least, of the feedback mechanisms are, therefore, stochastic and non-

<sup>4</sup> (UK) Engineering and Physical Sciences Research Council.

deterministic. Furthermore the system being modelled includes elements that contain implicit models of themselves. This is of course mathematically intractable posing a fundamental ultimate obstacle to the investigation [göd31, leh85]. Note also that convincing support for the hypothesis requires that the analysis and its associated predictive models, must necessarily be quantitative. The number of data points available for each of the classes of data is, however, likely to be relatively small. Statistical analysis and the determination of significance is, therefore, not straight forward. Work in the application of control theory to economic modelling [bec94] and, more recently, in the application of system dynamics to aspects of the software process [abd91, wae94, mad96] suggest, however, that progress is possible. Note also that software engineers and others in the organisations in which and with which they work do not, in general, have the understanding, knowledge, skills or experience required for this analysis as part of their background. The long term study requires, therefore, a collaborative, extended, multidisciplinary approach to achieve exploitable results.

The remainder of this paper focuses on the initial results of the first FEAST/1 black box study. Much still remains to be done and the most significant contribution of the present project may well be to arouse wider international interest and so trigger the necessary collaborative investigations.

## 4 A first case study: the Logica FW system

### 4.1 FW data

After extensive discussion with the collaborators and others, selection criteria and candidate systems have been identified. The preference was for medium to large systems, how ever defined. It was also considered desirable to concentrate on systems that were being used in a number of locations so that the effect of users' feedback which, it is believed, is likely to have significant impact, could be identified and assessed. Other criteria included the availability of historical data on system evolution to permit initial black box analysis to detect the presence or absence of feedback-like behaviour. Prior experience suggests that data on some ten releases is necessary for the identification of behaviour patterns. For the white box studies seeking to model internal process structure and to identify active feedback controls and their impact, ongoing projects were considered essential. Projects having, in addition, a sufficiently long history to provide meaningful black box data would be particularly useful since this would provide opportunities for linking characteristics inferred from the black and white box studies. At the time of writing collaborator products and processes satisfying these criteria have been identified, information on process structure and content is being gathered and metric data

should be available shortly.

The first system evolution data to be made available to the FEAST/1 project was on the Logica plc Fastwire (FW) financial transaction system. This 8 years old system is now installed on some one hundred sites. The data set received covers the most recent 5 years of its evolution. Since then there have been several main releases and many more sub-releases. The data set as received from Logica related to some 100 releases (as defined by them) with each entry including three data items; release ID, size in modules and number of modules changed. Release dates were also available for most of the data points. Many of these releases were, however, of the same size as their predecessor. Clarification revealed that these were fix releases that were very frequently only transmitted to those (limited) number of customers adversely affected by a fault in an earlier release. A subset of the data was therefore selected. As more familiarity with the system history has been attained the criteria, and therefore the subset selected, have had to be changed to yield the set shown in table 2. The analysis and plots presented below may, therefore, differ somewhat from those included in earlier publications [tur96], [leh96a,97]. The trends and patterns they display have, however, not changed significantly. Details of the selection and refinement criteria applied to the data have been documented in an internal report.

RSN	Size in Modules	Release ID	RSN	Size in Modules	Release ID
1	977	1.0	12	2087	5.0A
2	1344	2.0A	13	2091	5.0B
3	1390	2.0B	14	2095	5.0C
4	1492	2.0C	15	2101	5.0D
5	1581	2.0D	16	2151	5.0E
6	1595	2.0E	17	2167	5.0F
7	1800	3.0A	18	2312	6.0A
8	1832	3.0B	19	2315	6.0B
9	1897	4.0A	20	2696	7.0A
10	1897	4.0B	21	2699	7.0B
11	1902	4.0C			

Table 2 The FW data set


To protect their identity, the IDs of the releases listed in table 2 have been replaced by a sequence of identifiers that replace those assigned by Logica. In addition, and as was done in the OS/360 study [leh80b], consecutive integers have been assigned to the releases comprising the evolution sequence to be analysed. These provide a *pseudo-time* measure designated the *rsn*, a *sequence number* in the sense of Cox and Lewis [cox66]. Basing the analysis on this measure is appropriate because only at the instant of release of an *E*-type software system are its properties, as determined by the then established software

text, uniquely defined. By definition, an *E*-type system operates in a domain always liable to change at a rate that is accelerated by development, installation and operation of the system. Thus the software too, that is the code and/or its documentation, must be repeatedly updated and adapted<sup>5</sup> to remain a faithful model of the application in its operational domain. At the time of release the text is, by edict, fully defined. At all other times it is likely to be in a state of flux [leh85].

The releases included in the analysis whose results are presented below may be categorised into three classes:

- Major mainstream releases. These are intended to be adopted by the majority of user organisations. They are often required to achieve standardisation or for legal reasons.
- Minor mainstream releases. These provide minor improvements or enhancements. Such releases are included in the analysis presented below only where, in addition to other criteria, at least one module has been added or deleted with respect to its evolutionary predecessor.
- *Error correction* releases. These neither add nor enhance functionality. These also have been included in the analysis if they involve system growth by, at least, one module.

One consequence of ordering releases by rsn in the presence of the various types of releases is that a situation may arise in which the ordering adopted differs from the date ordering. For example, work on a new mainstream version 3.0A may have been proceeding concurrently with minor enhancement or bug-fixing of an older release. Release 3.0A may be shipped to mainstream clients before 2.0E is ready for delivery. As illustrated in figure 2, release 2.0D precedes release 3.0A in real time and might, therefore, appear to be its unique parent. Release 3.0A will, however, have also inherited functionality and code first developed for and integrated into 2.0E. Therefore, in the evolutionary sense release 2.0E is (at least) also a predecessor of 3.0A and is given the lower rsn.



**Fig. 2 Example of release ordering by date (not to scale)**

An additional release type has been identified in the FW process. This type, termed *ad-hoc*, is initially aimed at the satisfaction of the needs of a specific client. Such releases are excluded from the results presented below. However, unless providing some temporary facility later to be removed, the enhancements included in such releases are sooner or later integrated into the main stream to maintain smooth, uniform evolution over all installed systems and simplify overall FW configuration management. Moreover, these releases absorb project resources and, therefore, impact other concurrent activities. Thus they need, ultimately, to be included in the analysis.

## 4.2 System growth

With the cost of storage declining at all levels, system size is, in itself, not of major concern. It may, therefore, be seen as a independent and composite monitor of system evolution which, within limits, is neither planned nor managed. It is determined by other factors. Some of these will be managed, others "just happen". Size determinants include system design, programmer style and experience, development timetables and constraints, intensity of the desire to achieve compactness or clarity<sup>6</sup>. The great majority of reported software metrics work has tended to use *locs* (lines of code) as the measure of system size. As in the case of the original OS/360 study [leh80b,85] the FW analysis reported here has used module count for that purpose. In the absence of a better measure, module count also serves as an initial estimator of system functionality and power. The 1970s study did, in fact, compare the results of loc and module based studies. It was shown that these were essentially similar but with the locs measure providing a less consistent picture of the evolutionary behaviour of the system than did the module count. Locs are, therefore, considered inferior as a measure. The superiority of module count was explained by the observation that, however modules are defined, they have, within a given domain, some degree of functional integrity whereas locs have none [leh85]. The number of modules in a specific system is also not, in general, dependent on individual programmer practice. Module numbers may, therefore, be expected to provide a more consistent measure of system size and hence a better, though admittedly coarse, indicator of system functionality.

The *function point* (FP) [alb79] measure may be considered an alternative measure for system functionality or power. Their use does, however, raise some questions. For example, how arbitrary are the interpretation of the definitions or the factor ratings achieved and, as a consequence, how consistent are the results obtained by different raters? Moreover, the establishment of the measure requires judgment based on subjective measures and the overall determination is labour intensive and difficult to automate [kem93]. There is also little, if any,

<sup>5</sup> As per the first law of software evolution as reproduced in table 1.

<sup>6</sup> Or even, where productivity is measured in locs and the concern is with *productivity* improvement.

experience in application of the measure to larger systems. Finally, it must be observed that, unlike module count, FP data is not widely available from data archives of software systems across the industry. They are certainly not available from the systems being offered for study by the FEAST/1 collaborators. Module count is, therefore, being used as a size measure and, by implication, as a system power estimator, in the FEAST investigation. To date, and as illustrated by the results presented below, this decision appears to be justified.

The growth in modules of FW over releases  $rsn_1$  to  $rsn_{21}$ , that is releases 1.0 to 7.0B is shown in figure 3. The overall growth pattern should be compared with that of OS/360 over its first 20 or so releases as illustrated in figure 1.

The abscissa of figure 3 represents the individual release sequence numbers as explained above. The figure clearly shows the upward trend of system growth. The trend is, therefore, consistent with the first and sixth laws of software evolution but does not distinguish between them. Moreover, it also shows a ripple effect strongly reminiscent of that of OS/360 as in figure 1. It was, of course, this ripple phenomenon that first suggested that the software process was stabilised by feedback control, as captured in the third and eight laws. Thus this initial result of the present study is certainly compatible with the conclusions and, in particular, the laws of software evolution, first reached in the study of OS/360 more than twenty years ago.

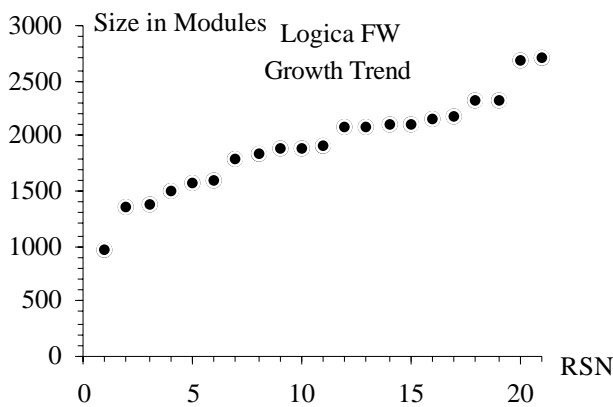


Fig. 3 FW growth trend by rsn

### 4.3 Incremental growth

Figures 4 and 5 show the incremental growth per release of OS/360 and FW respectively over the releases  $rsn_1$  to  $rsn_{21}$  for each system. The horizontal line indicates the average growth per release over this range. For FW the plot includes all the data in table 1. For OS/360 the final five releases for which data is available are omitted from the plot since they reflect the transition (in growth trend terms) from OS/370 to VS1 and VS2.

The two plots display remarkably similar cyclic characteristics though, as one should expect, they differ in detail. They also resemble statistical process control charts [dav84]. It was, in part, the observation of this cyclic pattern and its symmetry around the average with respect to OS/360 that originally led to formulation of the third and fifth laws. The long term trend of the moving average of the incremental growth of *E*-type systems as they evolve will, in general, be difficult to determine because of the small number of data points generally available. It might be the case that this average declines because of increasing complexity. Alternatively, it might grow as a consequence of improving process technology. It may also be that these (and other?) contrary pressures compensate for each other over time so that the original assertion of invariance remains valid. It is, in fact, not certain that all systems or domains behave in the same way. The fifth law as stated in table 1 will have to be re-examined. The analysis outlined below will provide additional insight for FW and will identify the growth trend model which, for that system at least, is to be preferred.

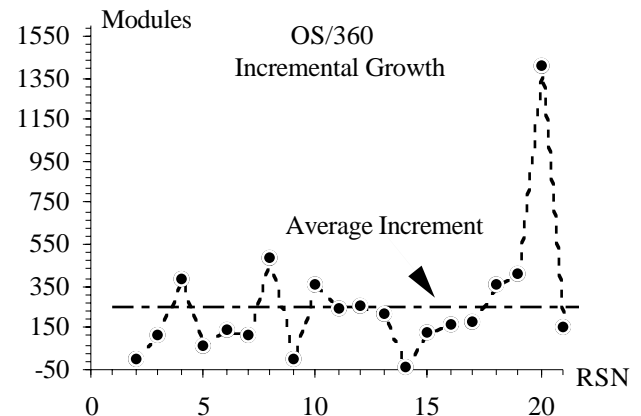


Fig. 4 Incremental growth per release of OS/360 over releases  $rsn_2$  to  $rsn_{21}$

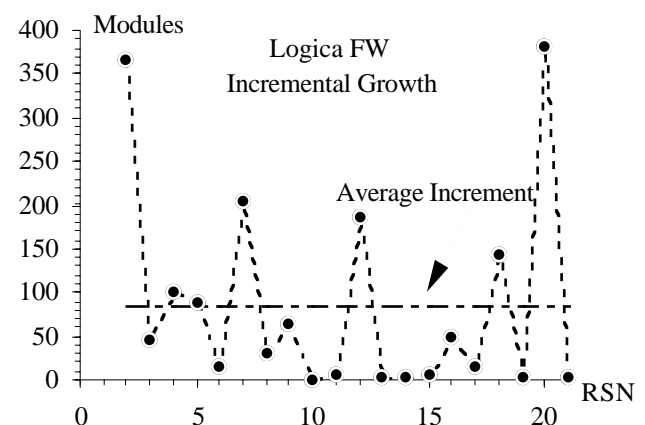


Fig. 5 Incremental growth per release of FW over  $rsn_2$  to  $rsn_{21}$

As pointed out previously [leh74,78], the cyclic effect reflected by the *peaks* and *troughs* in the incremental growth plots may be indicating the presence of feedback driven and controlled growth. Thus, influences tending to increase system functionality, that is growth towards the peaks, may have their source in positive feedback. The declines may reflect size stabilisation and other negative feedback effects. An example of such feedback is the evolutionary pressure that arises when clients and users express a need for enhancements to existing capability or system extension. But as implementation of such changes proceeds, the size and complexity of the system increases leading to declining comprehensibility, increasing error rates, increasing resistance to change or the impact of budgetary constraints. These lead to a decrease of resources available for, for example, growth as the resource demand for *fault fixing* and *complexity reduction* increases [leh85]. If sufficiently mature [hum95], the process will be directed in its evolution and growth patterns by data reflecting such needs. That is, the data or its derivatives will be used to adjust process objectives (immediate and/or long term) and process parameters. It will be used to drive, constrain, and in general, manage the process. Positive feedback drives growth while negative influences force a period of consolidation (correction and restructuring). An example of the consequences of excessive positive feedback may be provided by the final 7 releases,  $rsn_{20}$  to  $rsn_{26}$ , of OS/360 (figure 1). A hypothesis that explains the system's apparently unstable behaviour over these releases is that it was a consequence of excessive growth, in response to market demand, in going from  $rsn_{19}$  to  $rsn_{20}$ .

This brief analysis suggests that the FW data supports, in part at least, the third and fifth laws of software evolution as originally inferred from OS/360 study. Analysis of the long term growth trend of FW in the next subsection suggests, however, that the wording of laws III and V as in table 1, must be modified.

#### 4.4 The Inverse Square model (IS)

This section presents two models of FW growth. The first of these, illustrated in figure 6, is obtained from the data set of table 2 using a least squares linear (LSL) fit. The models focus on the general trend and largely ignore the ripple. Detailed analysis of the latter is beyond the scope of this paper.

After investigating other possibilities Turski developed an alternative, inverse square, model (IS) represented by the nonlinear discrete-time dynamical recursion (1) [tur96]. In this model  $s_i$  is the actual value of  $rsn_i$ ,  $\hat{s}_i$  is its fitted or predicted size, "n" is the total number of releases in the data set and  $\underline{E}$  is a model parameter.

$$\hat{s}_1 = s_1 \quad (1a)$$

$$\hat{s}_i = \hat{s}_{i-1} + \underline{E}/(\hat{s}_{i-1})^2 \quad \{i = 2, \dots, n\} \quad (1b)$$

The parameter  $\underline{E}$  is the average of individual  $E_i$ , calculated from either (2) or (3).

$$E_i = (s_i - s_{i-1})s_{i-1}^2 \quad \{i = 2, \dots, n\} \quad (2)$$

$$E_i = (s_i - s_1)/(\sum_{k=1}^{i-1} 1/(s_k)^2) \quad \{i = 2, \dots, n\} \quad (3)$$

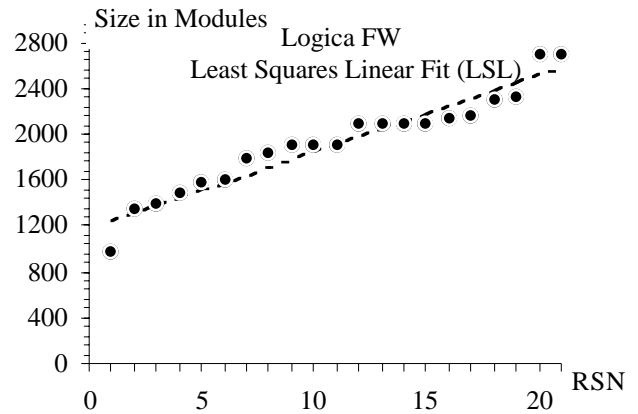


Fig. 6 Least squares linear fit to FW over  $rsn_1$  to  $rsn_{21}$

Algorithm (2) uses only the two most recent data points in computing  $E_i$ . With (3) all data to  $rsn_i$  are considered. In either case the average of the resultant set of  $E_i$  gives an estimated value for  $\underline{E}$ . A third approach (LSIS) computes  $\underline{E}$  from the entirety of data using a least squares criterion and is illustrated in figure 7.

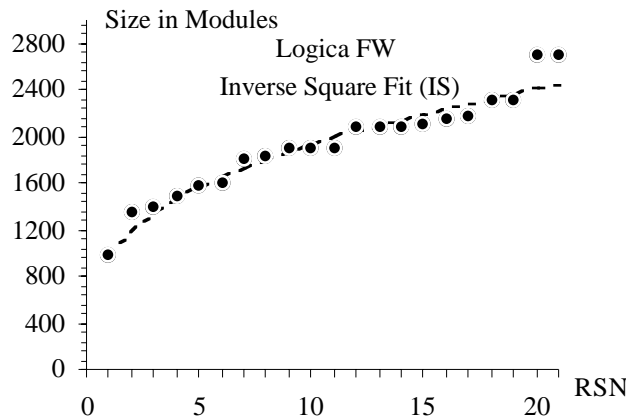


Fig. 7 Least squares inverse square fit to FW over  $rsn_1$  to  $rsn_{21}$

The conceptual implications guiding the selection of one of the three alternative algorithms for computing  $\underline{E}$  are subtle and are not discussed further here. They yield slightly different values for  $\underline{E}$  but, in the context of this study, they do not produce significantly different behavioural patterns. Nor do they change the conclusions to be drawn. Finally, the observant reader will notice apparent outliers  $rsn_{20}$  and  $rsn_{21}$ . No comment can be made at this time about the significance of these or their possible implication.

For the trend models estimated from the full set of 21 data points, statistical measures of the closeness of fit of the LSL and IS models do not differ significantly. Comparative assessment is, therefore, difficult on basis of currently available data. This may be due to the fact that the damping in the IS trend is not strong. Moreover, neither model addresses the ripple. The deviations from smooth growth that the latter represents could, of course simply be noise, the compounded impact of many, continuing, localised, often short term management and implementation decisions in which case it would not affect the assessment. The FEAST hypothesis suggested that, in part at least, the ripple is an indicator of the presence of feedback-controlled mechanisms that regulate the long term growth trend. The ongoing white box modelling activity in FEAST/1 represents a first step in the attempts to resolve this issue, to permit refinement of the models and a more precise assessment of the degree to which they, their derivatives or different models reflect the reality of the processes studied; and the degree to which they may be generalised.

Including the ripple will assist in comparative assessment of the model. It has, however, been pointed out already [tur96] that the phenomenology of the situation suggests several reasons for preferring the IS model:

- The IS inverse square property can be interpreted as reflecting the complexity growth of a software system over a sequence of releases. Such growth is due, in part, to increases in the complexity of the application, for example, as features not included in the original system definition, and often orthogonal to it, are added. Moreover, the process of evolution adds change upon change upon change with, in general, little attention paid to the resultant complexity growth [leh85]. It is this phenomenon that is captured by the second law (table 1).
- As a one parameter model IS is also compatible with the fourth law of software evolution, with the parameter  $\underline{E}$  reflecting the constant effort that the law identifies [leh78].
- IS also satisfies the Principle of Parsimony [cox66].
- No system can grow forever. The linear growth model is thus incompatible with reason and common experience.

#### 4.5 Further consideration of the Inverse Square model

The list of reasons for favouring IS over the LSL includes the observation that the single parameter  $\underline{E}$  of the former may be interpreted as a constant effort parameter as predicted by the fourth law. Estimation of  $\underline{E}$  from the available data strengthens that argument. Such estimation produces a value that, as shown below, remains relatively constant as FW evolves. That is, the single parameter of the model may be interpreted as the constant effort or work

output identified in the fourth law as being required to take the system from one release to the next. The principal questions raised by this interpretation, questions not satisfactorily answered, relate to the interpretation of  $\underline{E}$  and the units in which it is measured. Does  $\underline{E}$  relate to the input effort required to achieve release by release system evolution or to the output achieved from the process measured by some measure of increase in system quality and power? To answer the first question requires further investigation and additional data. As to units,  $s_i$  is a dimensionless count. Hence  $\underline{E}$  is dimensionless. But despite these unsolved questions it is concluded that, on the basis of currently available data, the above remarks, together provide some justification for preferring the IS model. It appears to reflect reality more closely.

The full implications of one further indicator of the superiority of IS over LSL must now be considered. When modelling large data sets, the first part is often used to estimate model parameters and the second to then evaluate its "predictive" capability [ger93]. With the small size of the data set available from FW, this might not appear to be a fruitful path to follow. Turski [tur96] did, however, investigate this question, asking: "How many points beginning with  $rsn_1$  have to be considered in order to get an appropriately low error of fit, an acceptable predictive capability?" In terms of the FEAST hypothesis this question is equivalent to asking: How fast is the FW dynamics established? An answer for FW is suggested by the plot of figure 8.


Figure 8 plots a set of mean absolute error of fit values ( $mae_j$  { $j = 2, \dots, 21$ }), where  $j$  indicates the number of points from  $rsn_1$  used to compute  $\underline{E}$ , see Appendix). The values of  $mae_2$  and  $mae_3$  are relatively large. As  $j$  is increased  $mae_j$  converges rapidly and reaches a relatively steady value by  $j$  equal 6 (parameter  $\underline{E}$  computed from the first six releases only). Thereafter  $mae_j$  { $j = 6, \dots, 21$ } has a mean of 74.6 with a standard deviation of 2.8. The  $mae_6$  value is only 4.7% of the system size at  $rsn_6$ , 3.2% of its size at  $rsn_{10}$  and 2.8% of its size at  $rsn_{21}$ . This behaviour is counter-intuitive in several ways. Possible interpretations and implications are summarised below. Overall, it does, however, appear to indicate the strength of the system dynamics. This phenomenon supports the observation made by one of the authors many years ago with regards to OS/360 evolution that "Rather than the managers managing the (evolving software) system, the system manages the managers." It must, of course, be understood that the reference here is to long term evolution, not to the specifics of individual decisions, often localised in time, system space and implementation space.

- Figure 8 based on the IS model suggests that the FW growth trend is established over some six of the releases included in the study. In accordance with the FEAST hypothesis, it is assumed that the dynamics arises from the characteristics of the software, the organisations developing, marketing and using the software, the communications between them and the

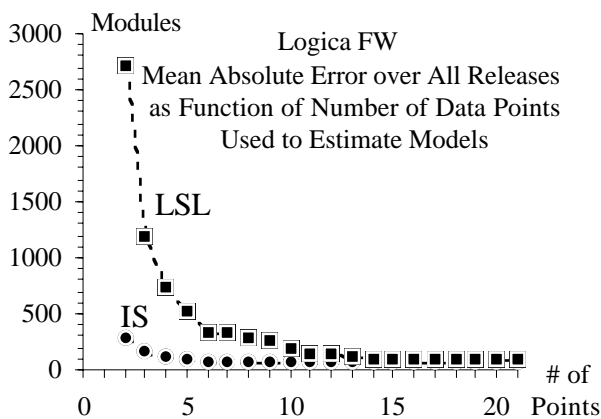


controls that are exercised. In any event figure 8 supports the hypothesis that the *E*-type systems evolution process develops strong dynamics.

- The *mae* for IS of 74.6 modules with standard deviation of 2.8 over the stable range is very close to the calculated average incremental growth of about 86.1 modules over all data points (fig. 5). This raises the question whether there is some relationship between the variance of the ripple (which is a significant source of error for the trend fit) and the mean incremental growth. Establishing a correlation would lead to a concept of *safe growth rate limits*. Establishing either would provide strong conceptual support for the *incremental or evolutionary release strategy* [gil88]. The entire question remains to be investigated.



**Fig. 8 Mean absolute error of fit to FW over all releases as function of number of points used to estimate IS model**



**Fig. 9 Mean absolute error of fit to FW over all releases as function of number of points used to estimate LSL model (squares) superimposed on that of IS model (circles)**

- Note that the *mae* of LSL over the stable range is, at 86 modules, even closer to the average incremental growth of 86.1 modules than is that of IS. The implications of this, for example on the evaluation of the relative value of the two models requires more investigation.
- The IS plot in Figs. 8 and 9 stabilises much more rapidly than does the LSL plot. Moreover, if IS and LSL are estimated by using only  $rsn_1$  and  $rsn_2$ , the former outperforms the latter by an order of magnitude. Thus while there are still unanswered questions, figures 8 and 9 appear to support the earlier conclusion that IS is to be preferred over LSL. That they provide further support for the FEAST hypothesis and the laws of software evolution does not require further emphasis

The results presented above are based on the examination of the FW system, investigation of OS/360 not having yet been reopened. Continued investigation of these and other systems is clearly required.

#### 4.6 Impact of the study on the laws of software evolution

More work is clearly required for firm conclusions to be reached in regards to the many issues raised above. It is nevertheless considered appropriate to indicate in table 3 the extent to which the investigators feel encouraged to see the present results as being compatible with, or even supporting, the laws of software evolution. The weight of evidence suggests that, despite the 20 year gap and the significant difference between IBM and Logica systems and their development and operational environments, there are strong similarities in the phenomenology of their evolutionary growth. It is believed that the results of the studies to date will, with some modification, extend to *E*-type systems in general. The FEAST/1 project will, it is hoped, receive sufficient data from the evolution processes of a variety of systems to establish confidence in a set of conclusions that are valid in some stated domain or, of course, to demonstrate that they cannot be generalised.

#### 5 Final remarks

The results achieved so far by applying this method in the FEAST/1 project are encouraging. Additional data on a wide spectrum of software systems to be received from various industrial collaborators should, if consistent, permit generalisation of both the conclusions reached and the measurement and analysis techniques being employed. The present paper describes the black box approach that has revealed aspects of FW evolution and of its evolution dynamics, has provided material for interpretation and for the formulation of explanatory hypotheses. A white box modelling approach is simultaneously seeking to model the structure of other industrial software processes and to simulate their behaviour including their feedback control

loops. These investigations are being further backed up through the development of a multi-agent model. It is hoped that this work will confirm, perhaps modified versions of, the laws of software evolution [leh96d] that now include the FEAST hypothesis and, put them on firmer foundations. If successful over a range of systems, the investigation will provide a base for a plausible theory of software process and software evolution. The alternative, that the results of the investigation demonstrate that the laws and the hypothesis are not of general relevance though satisfied for particular instances of *E*-type systems and their evolution processes cannot, at this stage, be dismissed.

The FEAST/1 study has already made visible progress in illustrating how measurement concepts can be applied to the study of software evolution. It has successfully extended the 1970s techniques by applying more rigour [law82] to mastery of the observed phenomena. The specific results derived are of considerable interest, both in themselves and from a wider perspective. The long term significance of this paper is, however, more likely to be in the approach and techniques it presents. Being able to

detect, measure and control feedback phenomena and their impact is believed to be key to major advances in software process management and execution.

In view of the fact that this paper will be presented at the Metrics '97 symposium it is appropriate to comment on its focus on the FEAST hypothesis, the related FEAST/1 project and the absence of references to other relevant metrics work [fen96,ieeee94,kit82,96,vot95]. FEAST/1 is believed to exemplify an original metrics based approach to the study of the software process and software evolution. This approach has been consistently followed from the first primitive study of OS/360 in the late sixties and seventies [leh69,85] to the current investigation. The study was triggered by a general observation; the universal and persistent problems accompanying software development and maintenance, ie. software evolution. Following recognition of the problem as appropriate for research investigation [leh69,85] and receipt of appropriate data, first from OS/360 and more recently from Logica FW [leh96d], patterns and regularities in their evolution were revealed and modelled. Interpretation of the models led, in turn, to the generation

No.	Brief Name	Support	Indicator
I	Continuing Change	√	Fig. 3 clearly indicates continuing growth. Logica's confirmation that this is partly due to adaptation and change supports the law. Quantification will be of interest.
II	Increasing Complexity	√	The inverse square law of growth (eq. 1) and its predictive power (fig. 7) supports complexity as a constraining factor.
III	Self Regulation	?	The ripple (fig. 3) of the, otherwise, smooth growth (eq. 1) suggests regulation around a smooth trend. Identification of the underlying mechanisms is required to support the law as it stands.
IV	Conservation of Organisational Stability (invariant work rate)	√	The ability to obtain a close fit and very good predictive power with a single and constant parameter $\underline{E}$ (eq. 1) provides support. Measures of the work rate are required.
V	Conservation of Familiarity	? <sup>7</sup>	Fig. 5 still suggests that the average incremental growth has a definite trend. Its invariance as in the original formulation is now, however, questioned. Determination of the trend and the consequences of a release whose incremental growth exceeds the average significantly must await the further behaviour of the system in its evolution.
VI	Continuing Growth	√	Fig. 3 clearly indicates continuing growth. Logica's confirmation that this is partly due to functional growth supports the law. Quantification will be of interest.
VII	Declining Quality	?	No data that provides evidence for or against is available.
VIII	Feedback System	√	Regulation as in figs. 3, 5, 7, 8 and inverse square law, (eq. 1) are supportive. Feedback control mechanisms must be identified to obtain further support.

**Table 3 The laws of software evolution in the light of the preliminary FW analysis<sup>7</sup>**

<sup>7</sup> It is hoped to obtain more data that will provide evidence, one way or the other.

of hypotheses (eg. the laws and FEAST) to interpret them. These successive steps led to an iterative investigation that is now yielding further data (historical and/or obtained by experimentation and measurement) to support, refute or modify and then to extend and generalise the emerging theoretical base and framework. Such results must, of course, be continually validated or adjusted by observation of and experimentation in actual industrial processes. Thus the more general relevance of the paper to the metrics community is in its approach which may be compared with those more widely adopted.

Apart from any theoretical advance that this study will provide, it should, if successful, lead to the development of methods and tools for process management, release planning and process improvement. This will shape the direction of software metrics, software process modelling and process improvement in the years to come. If the extent to which feedback phenomena in E-type evolution processes shapes and constrains the software process significantly, mastery and command of that phenomena will open up important new prospects. Moreover, the software process is a special case of business processes, in general [leh97]. The approach applied and the conclusions reached should find much wider application. It is believed that FEAST/1 is a study which, if successful, will eventually lead to a theory and to a technology which together can trigger major advances in the software and other business processes and their improvement.

## 6 Acknowledgements

We are grateful to Logica plc for providing access to the FW data and in particular to Joe Halberstadt for his collaboration. Sincere appreciation is also due to Profs. Berc Rustem and Vic Stenning, co-Principle Investigators on the FEAST/1 project, for their many contributions to the investigation and to Dr. Emma McCoy of the ICSTM Mathematics Department for her help with statistical aspects of this investigation. We also acknowledge the constructive contributions of participants in the three open FEAST workshops in 1994/5. Last but not least our thanks to the anonymous referees for their careful reading and constructive comments. Since October 1996 the work reported here has been supported under EPSRC grants numbers GR/K86008 and GR/L07437.

## 7 References<sup>8</sup>

- [abd91] Abdel-Hamid T and Madnick SE, *Software Project Dynamics - An Integrated Approach*, Prentice-Hall, Englewood Cliffs, 1991, 264 pps.
- [alb79] Albrecht AJ, *Measuring Application Development Productivity*, Proc. Guide/Share: IBM Application Development Symposium, Monterey, CA, 1979, pp. 83 - 92
- [bec94] Becker RS, Hall B, and Rustem E, *Robust Optimal Control of Stochastic Nonlinear Economic Systems*, J. of Economic Dynamics and Control, n. 18, 1994, pp. 125 - 148
- [bel72] Belady LA and Lehman MM, *An Introduction to Growth Dynamics*, Proc. Conf. on Statistical Comp. Perf. Evaluation, Brown Univ. 1971, Academic Press, 1972, W Freiberger (ed.), pp. 503 - 511
- [bel78]\* *id.*, *Characteristics of Large Systems*, Proc. Conf. Research Directions in Software Technology, (P. Wegner ed.), Sponsored by Tri-Services Committee of the DOD, Brown U. Providence, RI, Oct. 1878, MIT Press, 1979, pp. 106 - 142
- [box70] Box GP and Jenkins GM, *Time Series Analysis, Forecasting and Control*, Holden-Day, San Francisco, 1970, 553 pps.
- [cox66] Cox DR and Lewis PAW, *The Statistical Analysis of Series of Events*, Methuen, London, 1966
- [dav84] Davis OL and Goldsmith PL, *Statistical Methods in Research and Production*, 4th. ed., Longman, London, 1984, 478 pps.
- [fea94.5] *Preprints of the three FEAST Workshops*, Lehman MM (ed.), Dept. of Comp., ICSTM, 1994/5
- [fen96] Fenton NE and Pflieger SL, *Software Metrics - A Rigorous and Practical Approach*, 2nd ed., PWS Publ. Co., London, 1997, 638 pps.
- [for61] Forrester JW, *Industrial Dynamics*, Productivity Press, Cambridge, MA, 1961
- [for70] Forrester JW, *Understanding the Counter Intuitive Behaviour of Social Systems in Systems Behaviour*, ed. by Open Systems Group, 3rd. Ed., pp. 270-287, Paul Chapman Publishing Co. and The Open University, London, 1972
- [ger93] Gershenfeld NA and Weigend AS, *The Future of Time Series: Learning and Understanding*, in *Time Series Prediction: Forecasting the Future and Understanding the Past*, Gershenfeld NA and Weigend AS (eds.), SFI Studies in the Sciences of Complexity, Proc. Vol. XV, Addison-Wesley, 1993, pp. 1-70
- [gil88] Gilb T, *Principles of Software Engineering Management*, Addison Wesley, 1988
- [göd31] Gödel K, *Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme, I*, Monatshefte für Mathematik und Physik 38, 1931, pp. 173-1198. English translation, *On Formally Undecidable Propositions*, Gödel K, Basic Books, New York, 1962
- [hum95] Humphrey WS, *A Discipline for Software Engineering*, SEI Series in Software Engineering, Addison-Wesley, Reading, MA, 1995, 789 pps.
- [ieee94] *Measurement Based Process Improvement*, sp. iss. IEEE Softw., IEEE Comp. Soc. v. 11, n. 4, July 1994
- [kem93] Kemerer CF, *Reliability of Function Point Measurement: A Field Experiment*, CACM v. 3, n. 2, Feb. 1993, pp. 85 - 97
- [kit82] Kitchenham B, *System Evolution Dynamics of VME/B*, ICL Tech. J., May 1982, pp. 42 - 57
- [kit96] *id.*, *Software Metrics: Measurement for Software Process Improvement*, NCC Blackwell, 1996, 241 pps.

<sup>8</sup> Papers identified by a \* in the reference listing are reprinted in [leh85].

- [law82] Lawrence MJ, *An Examination of Evolution Dynamics*, Proc. 6th. Int. Conf. On Softw. Eng., Tokyo, Japan, 13 -16 Sept. 1982, IEEE Comp. Soc. Ord.N. 422, IEEE Cat n. 81CH1795-4, pp.188-196.
- [leh69]\* Lehman MM, *The Programming Process*, IBM Res. Rep. RC 2722, IBM Res. Centre, Yorktown Heights, NY 10594, Sept. 1969
- [leh74]\* *id.*, *Programs, Cities, Students, Limits to Growth?*, Inaugural Lecture, May 1974, Publ. in Imp. Col of Sc. Tech. Inaug. Lect. Ser., vol 9, 1970, 1974, pp. 211 - 229. Also in Programming Methodology, (Gries D , ed.), Springer, Verlag, 1978, pp. 42 - 62
- [leh77] Lehman MM and Patterson J, *Preliminary CCSS System Analysis Using Techniques of Evolution Dynamics*, Working Papers, First Software Life Cycle Management Workshop, Airlie VA, 1977, publ. by ISRAD/AIRMICS, Comp. Sys. Com., US Army, Fort Belvoir, VA, Dec. 1997, pp. 324 - 332
- [leh78]\* *id.*, *Laws of Program Evolution - Rules and Tools for Programming Management*, Proc. Infotech State of the Art Conf., Why Software Projects Fail, - Apr. 1978, pp. 11/1 11/25
- [leh80a]\* *id.*, *On Understanding Laws, Evolution and Conservation in the Large Program Life Cycle*, J. of Sys. and Software, v. 1, n. 3, 1980, pp. 213 - 221
- [leh80b]\* *id.*, *Programs, Life Cycles and Laws of Software Evolution*, Proc. IEEE Spec. Iss. on Softw. Eng., v. 68, n. 9, Sept. 1980, pp. 1060 - 1076
- [leh85]\* Lehman MM and Belady LA, *Program Evolution - Processes of Software Change*, Academic Press, London, 1985, 538 pps.
- [leh89] Lehman MM, *Uncertainty in Computer Application and its Control Through the Engineering of Software*, J. of Software Maintenance: Research and Practice, v. 1, n. 1, Sept. 1989, pp. 3 - 27
- [leh94] *id.*, *Feedback in the Software Evolution Process*, Keynote Addr., CSR Eleventh Annual Workshop on Softw. Evolution: Models and Metrics. Dublin, 7-9th Sept. 1994, in Information and Softw. Tech., sp. Iss. on Softw. Maintenance, v. 38, n. 11, 1996, Elsevier, 1996, pp. 681 - 686
- [leh96a] *id.*, *Process Improvement - The Way Forward*, Invited Keynote Address, Proc. Brazilian Softw. Eng. Conf., SBES'96, Universidade Federal de Sao Carlos, Brazil, 1996, pp. 23 - 35
- [leh96b] Lehman MM, Perry DE and Turski WM, *Why is it so Hard to Find Feedback Control in Software Processes?*, Inv. Pres., Proc. 19th Australasian Comp. Sc. Conf., Melbourne, Austr., Jan. 31 - Feb. 2, 1996
- [leh96c] Lehman MM and Stenning V, *FEAST/1: Case for Support*, ICSTM - DoC EPSRC Proposal, March 1996
- [leh96d] Lehman MM, *Laws of Software Evolution Revisited*, Position Paper, EWSPT96, Oct. 1996, LNCS 1149, Springer Verlag, 1997, pp. 108 - 124
- [leh97] *id.*, *Process Modelling - Where Next?*, ICSE 9 Most Influential Paper Award, Proc. ICSE 19, Boston, MA, 20 - 22 May 1997, pp. 549 - 552
- [mad96] Madachy RJ, *System Dynamics Modelling of an Inspection Process*, 18th Int. Conf. On Softw. Eng., Berlin, 25-29 March 1996, IEEE Comp. Soc. Ord. N. PR07246, pp. 376-386
- [mat92] *MATLAB High-performance Numeric Computation and Visualisation Software - Reference Guide*, The MathWorks, Inc., Natick, MA, 1992
- [mca95] McCabe FG and Clark KL, *Programming in April: An Agent Process Interaction Language*, in Intelligent Agents, Springer Verlag, 1995.
- [mcg93] McGowan CL and Bohner SA, *Model Based Process Assessments*, Proc. 15th Int. Conf. on Softw. Eng., Baltimore, MD, 17 - 21 May 1993, IEEE Comp. Soc. ord. n. 3700-02, pp. 202-211
- [mea72] Meadows DH et al, *Limits to Growth*, Signet, 1972
- [tur96] Turski W, Reference Model for Smooth Growth of Software Systems, IEEE Trans. on Softw. Eng., vol. 22, n. 8, 1996
- [vot95] Votta LG and Zajac ML, *A Design Process Improvement Case Study Using Process Waiver Data*, Proc. ESEC '95, Sitges, Barcelona, Spain, 25 - 28 Sept. 1995
- [wae94] Waeselynck H and Pfahl D, *System Dynamics Applied to the Modelling of Software Projects*, in Software - Concepts and Tools, v. 15, Springer-Verlag, Berlin, 1994, pp. 162 - 174
- [wil51] Wilkes M V, Wheeler D J and Gill S, *The Preparation of Programs for an Electronic Digital Computer*, Addison Wesley Press Inc., 1951, 167 pps.

## Appendix

This appendix indicates how the values of the mean average error of fit (*mae*), as in section 4.5 figures 8 and 9, have been computed. It also records the equations used to compute the least squares linear (LSL) and inverse square (IS) models.

As explained in section 4.5, for each of the models LSL and IS, a set of *mae<sub>j</sub>* values {*j*=2,...,21} was computed to determine the effect on the error of fit of the number of points used in the estimation. The average error over the entire data set for each such number of points was then taken as a measure of the goodness of fit for that model. Each set of *mae<sub>j</sub>* was calculated from the expression:

$$mae_j = (1/n) \sum_{k=1}^n |s_k - \hat{s}_{k,j}| \quad (A.1)$$

where throughout the appendix:

*n* (= 21 for the FW set) is the number of data points used in calculating *mae<sub>j</sub>*;

*j* {*j* = 2,...,21} represents the number of data points being used to estimate the LSL and IS models, respectively;

*s<sub>k</sub>* is the actual system size for the release with sequence number *rsn<sub>k</sub>* (table 2);

$\hat{s}_{k,j}$  represents the fitted system size for *rsn<sub>k</sub>*, with sub-index *j* indicating that the corresponding model (either LSL or IS) has been computed using the first *j* points of the data set only.

Similarly,  $\hat{s}_k$  is used below to represent the fitted system size based on LSL whose parameters have not, or IS whose parameter has not, necessarily been adjusted to minimise the error of fit in some sense.

### Computation of $\hat{s}_{k,j}$

**Least Squares Linear Model.** In this case  $\hat{s}_{k,j}$  is expressed as follows:

$$\hat{s}_{k,j} = a_j k + b_j \quad \{k = 1, \dots, n\} \quad (\text{A.2})$$

The parameters  $a_j$  and  $b_j$  are computed for each value of  $j$  using a least squares linear regression, as provided by most statistical packages and spreadsheets, to minimise  $\sum_{k=1}^j (s_k - \hat{s}_k)^2$ .

**Inverse Square Model.** For this model, each value of  $\hat{s}_{k,j}$  is computed recursively from  $s_1$  [tur96]:

$$\hat{s}_{1,j} = s_1 \quad (\text{A.3a})$$

$$\hat{s}_{k,j} = \hat{s}_{k-1,j} + E_j / (\hat{s}_{k-1,j})^2 \quad \{k = 2, \dots, n\} \quad (\text{A.3b})$$

where

$$E_j = (1/(j-1)) \sum_{i=2}^j E_i \quad \{j = 2, \dots, n\} \quad (\text{A.4})$$

$E_i$  in expression A.4 will have been computed either from expression 2 or 3 (section 4.4).  $E_j$  may also be computed using the least squares criterion and is then the value of  $E$  which minimises the error of fit,  $d_j$ , over the first  $j$  points, expressed as:

$$\min_E(d_j) = \min_E(\sum_{k=1}^j (s_k - \hat{s}_k)^2) \quad (\text{A.5})$$

where  $\min_E(\cdot)$  indicates the minimum value of  $(\cdot)$  over the entire range of the parameter  $E$ . Expression 1 (section 4.4) shows how  $\hat{s}_k$  may be computed.

For the FW data choosing one or other of these approaches has only minimally effect on the results. The choice has no significant impact on the interpretation of these results or on the conclusions reached.

Figures 7, 8 and 9 (IS plot) in section 4.5 are based on expression A.5. This has been implemented under MATLAB [mat92] and is available on the Web at <http://www-dse.doc.ic.ac.uk/~mml/feast1/>.