Integrated numeric and symbolic signal processing using a heterogeneous design environment

Ramamurthy Mani[†], S. Hamid Nawab[†], Joseph M. Winograd[†], Brian L. Evans[‡]

[†]Boston University
Electrical and Computer Engineering
Boston, MA 02215

[‡]University of California at Berkeley
Electrical Eng. and Computer Sciences
Berkeley, CA 94720-1772

ABSTRACT

We present a solution to a complex multi-tone transient detection problem to illustrate the integrated use of symbolic and numeric processing techniques which are supported by *well-established* underlying models. Examples of such models include synchronous dataflow for numeric processing and the blackboard paradigm for symbolic heuristic search. Our transient detection solution serves to emphasize the importance of developing system design methods and tools which can support the integrated use of well-established symbolic and numerical models of computation. Recently, we incorporated a blackboard-based model of computation underlying the Integrated Processing and Understanding of Signals (IPUS) paradigm into a system-level design environment for numeric processing called Ptolemy. Using the IPUS / Ptolemy environment, we are implementing our solution to the multi-tone transient detection problem.

Keywords: design environments, symbolic signal processing, heuristic search, blackboard model, synchronous dataflow model, Ptolemy, Integrated Processing and Understanding of Signals (IPUS), multi-tone transient detection, Adaptive Time-Frequency Representation (ATFR)

1 INTRODUCTION

Traditional solutions to complex signal processing problems may be viewed as consisting of (1) numerical algorithmic solutions to idealized mathematical formulations of various subproblems and (2) incorporation of symbolic heuristic search for the run-time selection of various parameter values (e.g. analysis interval boundaries, FFT length, filter length, detection threshold) associated with the algorithmic solutions. Software tools such as Matlab / Simulink, SPW, and Ptolemy are now commonly used to model, simulate and test the numerical processing aspects of signal processing solutions. Such software tools are based on the recognition that many numerical signal processing algorithms have underlying commonalities in their computational models (e.g. dataflow, discrete-event, and imperative) and data structures (e.g. complex numbers, vectors, and matrices). While computational structures such as the blackboard model[1] have been formulated and used for conducting symbolic heuristic search, it is desirable to have such structures available within the *same* software tools that also allow the user to utilize common numeric processing structures. This has recently been achieved through a joint effort between the Knowledge-Based Signal Processing Group at Boston University and the Ptolemy Group at the University of California at Berkeley. In particular, the ICP tool[2,3] for blackboard-based heuristic search has been incorporated as a domain within the heterogeneous design environment of Ptolemy.[4] In this paper, we describe a multi-tone transient detection problem for which we have formulated a solution whose implementation can take advantage of the incorporation of ICP within Ptolemy. This example serves two major purposes. First and foremost, it serves to illustrate the types of heuristic search that can be conveniently implemented within the ICP-augmented Ptolemy environment. Secondly, the example itself represents a novel solution to the problem of multi-tone transient detection.

The problem of detecting multiple and possibly simultaneous narrowband transients in the presence of other interfering signals arises in a variety of applications including music transcription, auditory scene analysis, and anti-jamming CDMA receivers for spread spectrum communication. Signal processing solutions involving time-frequency analysis, signal modeling, matched filtering and / or parameter estimation can be brought to bear on various aspects of this problem. However, for realistically complex scenarios, it is generally recognized that it is difficult to match the assumptions behind various signal processing solutions to the dynamically changing signal context for the detection problem. We have investigated the utilization of heuristic search methods to aid this process of matching signal processing solutions to the dynamically changing context. The flexibility offered by the incorporation of search processes also opens up the possibility of constructing mathematically tractable signal processing solutions for specialized contexts.

An example of a specialized context for transient detection is that of a single exponentially modulated narrowband signal immersed in noise. We have formulated a detector based upon a parametric exponential signal model for detecting amplitude transients. The design of the detector takes advantage of the stationarity of the value of an exponential rate parameter over the duration of the transient. Our theoretical analysis of the detector and experimental results with synthetic and real data have shown it to be very effective in its specialized context. However, we have to ensure that the detector is applied only when its underlying assumptions are satisfied. The traditional signal processing approach to accomplish this involves the use of a filterbank aimed at separating individual signal components. The key problem is then one of searching for the appropriate filters to be used in the filterbank. Furthermore, the search has to be performed online when the signal context is expected to keep changing over time. The formulation of a solution to this problem that incorporates heuristic search is one of our major results.

2 TRANSIENT DETECTION: PROBLEM DESCRIPTION

The problem of detecting the transient regions of each narrowband component in a signal with multiple narrowband components may be divided into three subproblems of increasing complexity. We now describe each of these subproblems and the main issues that need to be addressed while resolving them.

As illustrated in Figure 1(a), we first consider the problem of detecting the transient regions of a single fixed frequency narrowband signal contaminated by white Gaussian noise. We refer to this as the AM tone subproblem. We assume that the narrowband signal is of the form:

$$x(t) = a(t)\exp(j\omega_0 t + \phi), \tag{1}$$

where $a(t)$ is the amplitude modulation, ω_0 is the frequency, and ϕ is an arbitrary phase factor. The problem of detecting the transient regions is now one of finding time regions of $a(t)$ which correspond to signal onsets and offsets. In Section 3, we present a mathematical formulation to solve this problem. This solution does not require the use of heuristic search.

The problem of transient detection in an amplitude and frequency modulated tone immersed in white Gaus-

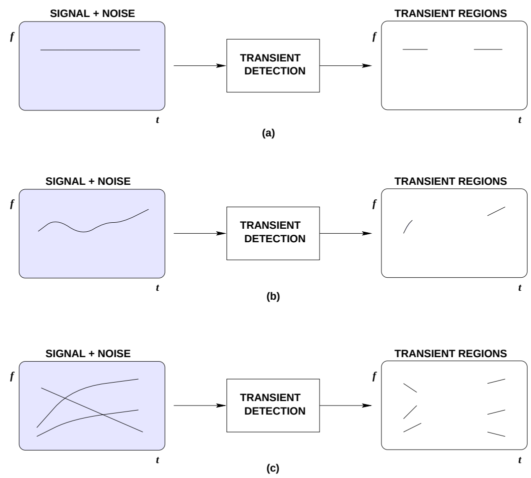

Figure 1: The three subproblems of the overall transient detection problem. (a) AM tone subproblem (b) AM/FM tone subproblem (c) Multi-tone subproblem

sian noise (the AM / FM tone subproblem illustrated in Figure 1(b)), involves signals of the form:

$$x(t) = a(t) \exp \left\{ j \left(\int_0^t \omega(\sigma) d\sigma + \phi \right) \right\}, \tag{2}$$

where $\omega(t)$ is the instantaneous frequency of $x(t)$. In Section 4 we describe how we have extended our AM tone solution to take into account the effects of frequency modulation. Once again, this solution does not require the use of heuristic search.

The need for heuristic search arises when we consider transient detection in the context of multi-tone signals. We refer to this as the multi-tone subproblem (see Figure 1(c)). Here, the signal consists of multiple amplitude and frequency modulated narrowband components:

$$x(t) = \sum_{k=1}^{N} a_k(t) \exp \left\{ j \left(\int_0^t \omega_k(\sigma) d\sigma + \phi_k \right) \right\}. \tag{3}$$

In order to address this problem, we have to first identify a technique for isolating individual narrowband components. This technique should be equipped to handle conflicting time and frequency resolution issues that are bound to arise within a multi-tone signal. Once individual components have been isolated, the problem of detecting transient regions may be addressed in a manner similar to that used for the AM and AM / FM tone subproblems. In Section 5 we present an approach for addressing the multi-tone subproblem. In particular, the approach makes use of heuristic search methods in conjunction with a time-frequency analysis of the multi-tone signal in order to isolate the various tones.

3 SOLUTION TO AM TONE SUBPROBLEM

Our solution to the AM tone subproblem makes use of a nonlinear estimator of the instantaneous exponential rates associated with a single tone signal. The exponential rate is defined using the following signal model:

$$x(t) = A \exp \left\{ \int_0^t (r(\tau) + jw(\tau))\, d\tau \right\}, \tag{4}$$

where $r(t)$ and $w(t)$ are real-valued functions of t. The function $w(t)$, which is the instantaneous frequency of $x(t)$, is a constant for tonal signals with no frequency modulation. The function $r(t)$ represents the instantaneous exponential rate of $x(t)$. We note that the function $r(t)$ is a constant for purely exponential transients. Furthermore, different instances of the same transient signal share the same $r(t)$ even if their amplitudes are scaled differently. This stationarity of the exponential rate allows us to build a transient detector based upon applying a simple threshold to exponential rate estimates.

3.1 Exponential transient detector (ETD)

An estimator[5] for the exponential rate is developed by noting that we may write (4) as

$$r(t) = \frac{d}{dt}\left(\ln|x(t)|\right) = g\{x(t)\} \tag{5}$$

We may thus view $r(t)$ as being obtained by applying a *continuous-time exponential rate operator (ERO)* $g\{.\}$ to the signal $x(t)$. If Δ is a short interval of time, then we may also express (5) as

$$r(t) = \frac{\ln|x(t)| - \ln|x(t-\Delta)|}{\Delta} = \frac{1}{\Delta}\ln\left(\frac{|x(t)|}{|x(t-\Delta)|}\right). \tag{6}$$

For facilitating the digital processing of signals, we have formulated a discrete-time approximation to our continuous-time rate operator. This approximation is obtained from (6) by replacing the time-decrement Δ with NT, where T is the sampling period and N is a positive integer. The discrete-time *ERO* is given by :

$$\mathcal{G}\left\{x[n]\right\} = \frac{1}{NT}\ln\left(\frac{|x[n]|}{|x[n-N]|}\right), \tag{7}$$

where $x[n] = x(nT)$.

The discrete-time ERO output has values in the range $(-\infty, \infty)$. For convenience, we have imposed a normalization on the ERO that performs a bijective mapping of values in $(-\infty, \infty)$ to values in $(-1, 1)$, preserving symmetry about 0. The normalized discrete-time ERO is given by:

$$G\{x([n]\} = \tanh\left(\frac{NT}{2}\mathcal{G}\left\{x[n]\right\}\right) = \frac{|x[n]| - |x[n-N]|}{|x[n]| + |x[n-N]|} \tag{8}$$

Decay transients have normalized instantaneous exponential rate values in the range $(-1, 0]$, while attack transients have values in the range $[0, 1)$. Furthermore, transient detection may be performed by imposing suitable thresholds on the normalized exponential rate. We refer to the resulting detector as the *Exponential Transient Detector (ETD).*

The advantage of using the ETD over using a detector based on applying a threshold to amplitude derivatives is illustrated through the following example involving noiseless exponential transients. For a particular assumed distribution of exponential rates, in Figure 2 we plot the probability of detection of a transient as a function of transient amplitude for both the ETD and the derivative-based detector. The probability of false alarm was assumed to be zero for both detectors. It is obvious from the figure that the ETD has the advantage that its performance does not degrade with decreasing signal amplitude.

3.2 ETD for noisy signals

We have developed a three-stage ETD for detecting the transient regions of a narrowband signal contaminated with white noise. While building this ETD, it is assumed that preliminary bandpass filtering has been

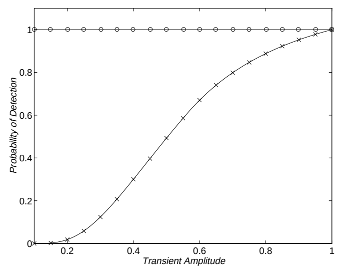

Figure 2: Advantage of ETD over a derivative-based detector. The line with $\circ$ shows a plot of P_D as a function of signal amplitude for the ETD. The line with $\times$ shows a plot of P_D as a function of signal amplitude for the derivative-based detector. Clearly, the ETD has the advantage that its performance does not degrade with decreasing signal amplitude.

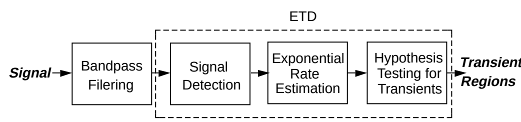

Figure 3: Approach for detecting the transient regions of a narrowband signal immersed in noise

applied on the narrowband signal to increase the effective signal-to-noise ratio. The three stage ETD is shown in Figure 3. In the first stage, signal detection is performed by applying a threshold on the magnitude of the bandpass filter output. The ERO is then utilized to estimate exponential rates in time regions where signal detections have occurred. Hypothesis testing for transient regions is performed by applying a threshold on the exponential rate estimates. This threshold is chosen such that the probability of detecting a non-transient region as a transient is kept below a specified bound.

4 SOLUTION TO AM/FM TONE SUBPROBLEM

Our solution to the AM tone subproblem can be extended to frequency modulated narrowband signals by preceding it with a time-varying bandpass filter[6] designed to track the center frequency of the tone (see Figure 4). At each time instant, the center frequency of the filter is initially set to a coarse estimate of the instantaneous frequency of the narrowband component of interest. The center frequency of the filter is then iteratively refined on the basis of progressively better estimates of the instantaneous frequency obtained from the filter's output. These estimates are obtained through the use of the energy separation algorithm[7] involving the Teager-Kaiser operator (TKO).[7] Using this algorithm, the instantaneous frequency of the bandpass filter output $x(t)$ is given by:

$$\omega(t) = \sqrt{\frac{\Psi[\dot{x}(t)]}{\Psi[x(t)]}} \tag{9}$$

where $\dot{x}(t) = \frac{dx(t)}{dt}$. Also,

$$\Psi[x(t)] = \left(\frac{dx(t)}{dt}\right)^2 - x(t)\left(\frac{d^2x(t)}{dt^2}\right) \quad \text{and} \quad \Psi[\dot{x}(t)] = \left(\frac{d^2x(t)}{dt^2}\right)^2 - \frac{dx(t)}{dt}\left(\frac{d^3x(t)}{dt^3}\right) \tag{10}$$

are the TKO outputs for $x(t)$ and $\dot{x}(t)$, respectively. The advantage of this iterative bandpass filtering approach is that the TKO-based algorithm used for performing the instantaneous frequency estimation is computationally very efficient.

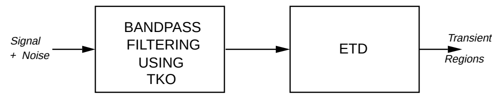

Figure 4: Approach for detecting transient regions of an AM/FM tone.

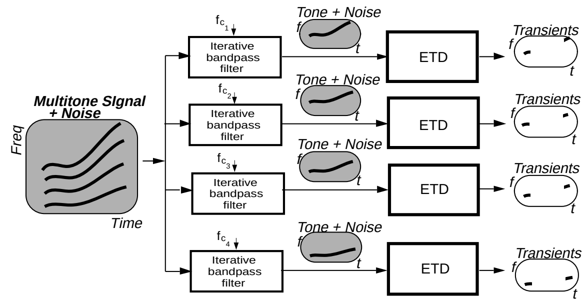

Figure 5: Transient detection in multi-tone signals.

5 SOLUTION TO MULTI-TONE SUBPROBLEM

An important component of our basic approach to the detection of transients in a multi-tone signal is to perform over each analysis interval of the signal a time-frequency analysis with the aim of identifying the separate AM/FM tones contained within it. This identification need not be overly precise in tracking the center frequencies and amplitudes of each of the tonal components. However, it is important for each tone to be isolated within a subregion of the time-frequency plane where there is little or no contribution from the other tones. For this *tone isolation* problem, we have formulated a heuristic-search solution that can be implemented in the ICP-augmented Ptolemy environment. Once tone isolation has been accomplished, the multi-tone signal is passed through a bank of iterative bandpass filters. Each iterative bandpass filter is selected to correspond to one of the isolated tones. The initial coarse estimates of center frequency used in each iterative bandpass filter are selected in accordance with the results from the time-frequency analysis for the corresponding tone. Furthermore, the bandwidth of each iterative bandpass filter is selected to ensure minimal contribution from neighboring tones. Exponential transient detection may then be applied to each of the isolated tones. This process of using a bank of iterative bandpass filters followed by exponential transient detection is illustrated in Figure 5. Let us now consider in detail the problem of tone isolation through time-frequency analysis.

5.1 Exhaustive search for tone isolation

A major issue in time-frequency analysis for tone isolation is the data-dependent optimization of conflicting time resolution and frequency resolution requirements. One approach for such optimization in the time-frequency plane has been reported by Parks and Jones.[8] Their approach, known as the adaptive time-frequency representation (ATFR), performs at each point in the time-frequency plane an exhaustive search within a large set of analysis windows in order to identify a window which optimizes a specific kurtosis criterion. While the ATFR approach has been shown to be effective for analyzing multi-tone signals, its major drawback is that it is computationally very expensive, typically requiring a search over $\mathcal{O}(10^2)$ windows at each time-frequency point. It should be noted that for each candidate window a short-time Fourier transform has to be computed over a time-frequency subregion of the time-frequency point under consideration.

Given the objective of tone isolation, it should generally not be necessary to optimize the time-frequency resolution tradeoff in all subregions of the time-frequency plane. This idea is illustrated in Figure 6. The gray areas indicate regions in which the application of ATFR may be necessary. On the basis of this idea, we have formulated an approach to tone isolation which uses symbolic heuristic search to restrict the application of the ATFR solution only to those regions where it is deemed necessary. This heuristic search approach to tone isolation is described in the next section.

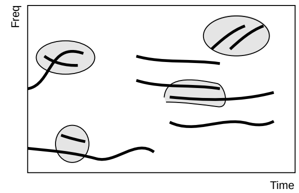

Figure 6: Figure illustrating the restriction of ATFR processing to a few regions in the time-frequency plane. Gray areas indicate regions in which the application of ATFR may be necessary.

5.2 Heuristic search for tone isolation

Based upon the IPUS blackboard paradigm[9] for the integrated processing and understanding of signals, our heuristic search approach makes use of *prediction* heuristics, *discrepancy-detection* heuristics and *reprocessing* heuristics. A key feature of this approach is that the processing of any particular signal interval is guided by the results obtained from the analysis of previously processed intervals. In particular, prediction heuristics are used to anticipate the likely evolution of tonal components from the previous analysis interval into the current one. These predictions form the basis for identifying and applying ATFR processing to the current interval with fixed analysis windows at each time-frequency location. We refer to such processing as searchless ATFR processing. The result of searchless ATFR processing is then compared to the original predictions using a set of discrepancy-detection heuristics. The identified discrepancies are then used in conjunction with reprocessing heuristics in order to select time-frequency subregions over which to apply the exhaustive search ATFR technique. We will now illustrate this approach through an example.

Let us consider a signal scenario which consists of two single-tone signal components as shown in Figure 7(a). Let us now assume that the signal has been processed until time $t = 1$ sec and this has resulted in the isolation of the two single tone signal components as shown by the solid lines in Figure 7(b). Furthermore, let us assume that we wish to apply the heuristic search for tone isolation over the current analysis interval from $t = 1$ sec until $t = 1.25$ sec. The processing results that are available until time $t = 1$ sec are used to perform a symbolic search for the appropriate prediction heuristic to utilize in the current analysis interval. This leads to the selection of the following prediction heuristic:

```
NAME:  Steady-Tone Prediction Heuristic
CONDITIONS FOR APPLICATION:
     A steady tone exists at time t = t0 sec, the start of an analysis interval.  (Here, a
     tone is defined to be steady at time t = t0 if its center frequency has not varied by
     more than ±5 % over the time range from t = t0 − 0.25 sec to t = t0 sec.)
ACTION:
     Post the hypothesis that the tone will continue over the time range from t = t0 sec
     until t = t0 + 0.25 sec with a maximum variation of ±5 % to its center frequency.
```

The utilization of this prediction heuristic over the current analysis interval from $t = 1$ sec until $t = 1.25$ sec results in the anticipation of tones at 650 ± 32.5 Hz and 900 ± 45 Hz. The anticipated locations of these tones are shown by the dotted lines in Figure 7(b). On the basis of these predicted locations for single-tone components, searchless ATFR processing is performed on the current analysis interval. The analysis window at each time-frequency point is chosen such that the searchless ATFR processing provides adequate frequency resolution for isolating tones at 682.5 Hz and 850 Hz (the closest possible frequency locations for the two predicted steady

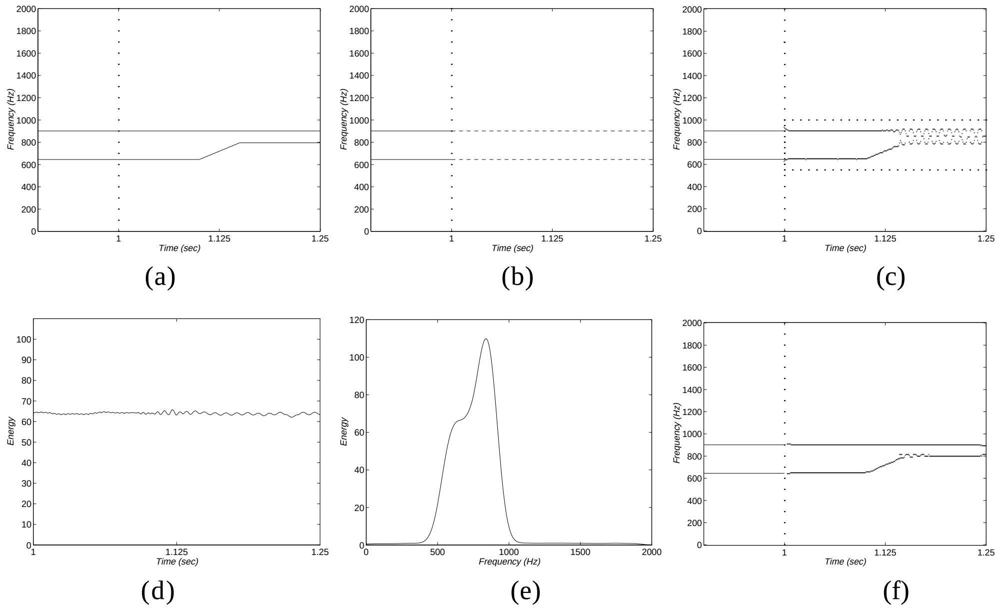

Figure 7: Example illustrating our heuristic search approach. Plot (a) shows the two tone signal scenario considered in this example. In plot (b) the black lines indicate the single-tone components that have been isolated as a result of processing the signal until time $t = 1$ sec. The dotted lines indicate the predicted locations for single-tone signal components in the analysis interval from $t = 1$ sec to $t = 1.25$ sec. Plot (c) shows the result of searchless ATFR processing. Plots (d) and (e) show the temporal and spectral energy profiles computed from the results of the searchless ATFR processing. Plot (e) shows the results of applying exhaustive search ATFR to the discrepant region shown within dotted lines in plot (c).

tones). Let us denote the magnitude-squared of the output of ATFR processing by $G(t, f)$, where t is the time variable with values in the range $1 \leq t \leq 1.25$ and f is the frequency variable with values in the range $0 \leq f \leq 2000$ (based on the assumption that the multi-tone signal considered here has negligible energy above 2000 Hz). The prominent peaks in $G(t, f)$ are plotted in Figure 7(c). In order to facilitate discrepancy-detection, $G(t, f)$ is also used to compute temporal and spectral energy profiles of the current analysis interval. While the spectral energy profile is computed by integrating $G(t, f)$ over t for $1 \leq t \leq 1.25$, the temporal energy profile is computed by integrating $G(t, f)$ over f for $0 \leq f \leq 2000$. The temporal energy profile and the spectral energy profile are plotted in Figures 7(d) and (e), respectively. The predictions that have been made over the current analysis interval are then used to perform a symbolic search for one or more suitable discrepancy-detection heuristics that must be utilized. This leads to the selection of the following discrepancy-detection heuristic as one of the heuristics that needs to be used over the current analysis interval:

NAME: Steady-Tone Discrepancy-Detection Heuristic
CONDITIONS FOR APPLICATION:
Applicable in a frequency subregion from $f = f_0$ Hz to $f = f_1$ Hz over which N steady tones have been predicted.
TEST:
Are there exactly N peaks in the spectral energy profile over the frequency subregion from $f = f_0$ Hz to $f = f_1$ Hz?
ACTION WHEN TEST FAILS:
Post a Peak-Insertion discrepancy hypothesis if there are more than N peaks.
Post a Peak-Deletion discrepancy hypothesis if there are less than N peaks.

In the current analysis interval the test applied in the above discrepancy-detection heuristic requires that there are exactly 2 peaks in the spectral energy profile in the frequency subregion from 617.5 Hz to 945 Hz (that is, the frequency subregion over which the two tones are predicted to lie). An inspection of Figure 7(e) shows that there is only one peak in this subregion, and therefore the test fails and a Peak-Deletion discrepancy is declared in the frequency subregion from 617.5 Hz to 945 Hz. This subregion is indicated by a dotted box in Figure 7(c). The identified discrepancy is now used in conjunction with the temporal energy profile in order to perform a search for the appropriate reprocessing heuristic to apply in the discrepant time-frequency regions. This leads to the selection of the following reprocessing heuristic:

```
NAME:   Steady-Tone Frequency-Shift Reprocessing Heuristic
CONDITIONS FOR APPLICATION:
     Applicable when a Peak-Deletion discrepancy has been identified and
     there is no change of more than ±20 % in the temporal energy
     profile of the analysis interval.
ACTION:
     Apply exhaustive ATFR search over the entire analysis interval within the
     frequency subregion over which the discrepancy exists.
```

This heuristic is selected on the basis that the absence of any abrupt changes in the temporal energy profile rules out the origination of a new single-tone component or the termination of a predicted single-tone component. Therefore, the discrepancy must be as a result of a significant shift in frequency of one or both tones. Upon utilizing the above heuristic, exhaustive search ATFR is applied to the time-frequency subregion over which the discrepancy was identified (the region within the dotted box in Figure 7(c)). The result of reprocessing this subregion is shown in Figure 7(f). We see from this result that the frequency location of the 650 Hz tone has changed to 800 Hz.

The above example provides a sampling of the heuristics that are employed in our heuristic-search approach for tone isolation. A number of other heuristics such as those related to prediction of chirping single-tone components and those related to detecting discrepancies resulting from the origination of new single-tone signal components have also been developed. During the course of implementing our solution to the multi-tone transient detection problem in the ICP-augmented Ptolemy environment, we are experimenting with a wide variety of heuristics with the aim of refining and enhancing the set of heuristics that we have already developed. The augmented Ptolemy environment is especially suited for carrying out such experimentation because it facilitates the implementation of these heuristics as modules which may be altered or substituted with ease. In the next section, we provide an overview of the main features of the ICP-augmented Ptolemy environment.

6 IPUS DOMAIN FOR PTOLEMY

Our approach in efficiently detecting multi-tone transients requires a combination of symbolic and numeric processing. The symbolic processing applies a blackboard architecture to determine the number of tones and to approximate their locations. The numeric processing uses this information to extract the transients by applying an iterative bandpass filter followed by an exponential transient detector to each tone. To simulate and implement our system, we need an environment that combines models for heuristic search and signal processing. In order to experiment with alternative heuristic strategies, we desire an interactive and modular design tool that supports iterative design.

Ad hoc techniques for symbolic processing have been incorporated in many signal processing systems. Currently available tools and languages used for signal processing have little formal support for mixing symbolic and numeric processing. Neither dataflow-oriented design tools such as SPW by the Alta Group of Cadence Design Systems Inc. nor implementation languages such as Fortran and Matlab[10] offer the sophisticated control structures for heuristic search.

We have incorporated the IPUS C++ Platform (ICP)[2,3] into the Ptolemy[4] system-level design environment. Ptolemy is an open architecture on which a collection of application-specific tools called domains are built. Domains can cooperate during simulation and, to a lesser extent, code generation. ICP has been encapsulated as an IPUS domain in Ptolemy. Through the IPUS domain, the ICP user gains access to a large library of signal processing functions, a variety of control and communication architectures, and an integrated graphical design environment. The Ptolemy user gains access to a knowledge-based architecture that can be used for sophisticated signal processing, signal reprocessing, and associated control strategies.

6.1 Background

The Ptolemy project advocates an important philosophy for tool development, that of the *heterogeneous design environment*. Instead of relying on a single approach, the project develops many different models of computation, control, and communication that can be combined hierarchically to form complex heterogeneous systems. Research ideas are implemented and tested in the Ptolemy software environment. The Ptolemy software environment provides a common framework to support the interaction of independent special-purpose tools for designing application-specific subsystems. This approach follows from recognition of the inverse relationship between adaptation and flexibility for design tools— that the more highly adapted an approach is to solving a particular problem, the worse it is likely perform when applied to others. The Ptolemy software environment is a freely distributable open software architecture with established strengths in mixing design tools for dataflow, discrete-event, and finite-state control systems, as well as tools for hardware / software cosimulation and automatic code generation. In several domains, Ptolemy provides not only the infrastructure for design and development, but also a series of heavily populated function libraries that the system designer can draw on for many standard system components. Written in C++, Ptolemy relies heavily upon the object-oriented design principles of polymorphism and encapsulation to obtain heterogeneity. These features have been relied on greatly in our extensions for the support of symbolic processing, as discussed below.

ICP was originally created to support the development of embedded applications employing the IPUS (Integrated Processing and Understanding of Signals) control architecture.[9] It provides a blackboard system to support a hierarchy of abstract signal representations[11] and a flexible and efficient planning mechanism[12] for performing run-time scheduling of signal processing actions based on application specific schema. ICP is comprised of a class library of system components and a set of extensions to the C++ language designed to facilitate the process of deriving application-specific instances of these objects. A pre-processor is used to convert these language extensions directly into C++ syntax, allowing efficient and embeddable applications to be obtained directly. While ICP provides the support for incorporating symbolic heuristic search into signal processing applications, it does not provide any functions for performing signal processing actions— these must be provided by the system designer. It also provides no direct support for other important system control architectures, such as those previously integrated with the Ptolemy environment.

6.2 The new IPUS domain

The IPUS domain for Ptolemy brings the complete functionality of ICP into the Ptolemy framework. Using the multiple inheritance feature of the C++ language, the IPUS domain kernel provides a set of generic system components that share the properties associated with both Ptolemy and ICP. Thus, existing ICP applications can be run within the IPUS domain, and those developed within the Ptolemy environment can also be run in stand-alone ICP, provided that any functionality obtained from other Ptolemy domains can also be migrated out of the environment (e.g. through code generation).

Domains cooperate with each other indirectly through the Ptolemy kernel by passing data and control back and forth to the kernel in a universal format. Integrating a new domain into Ptolemy only requires a translator of its internal data and control representations to and from this universal format. These translators have been implemented for the IPUS domain and tested for interactions between system components developed in the

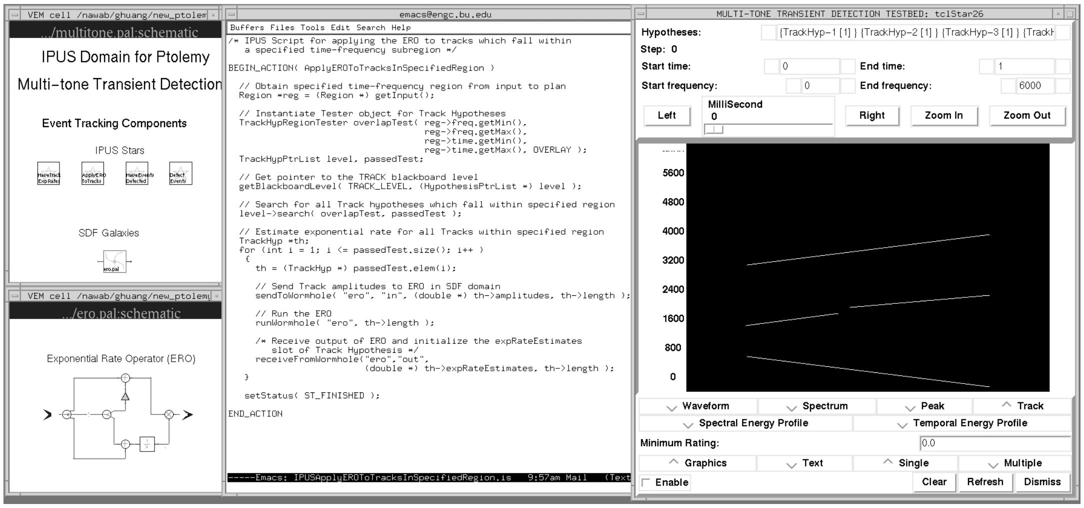

Figure 8: A screen image taken from the multi-tone transient detection application under development using the IPUS domain for Ptolemy.

IPUS and synchronous dataflow domains.

Several aspects of the IPUS domain are illustrated in Figure 8, which shows a screen image taken from the development of the multi-tone transient detection application. In the upper-left window of the figure, the Ptolemy environment's graphical representation of a portion of the application is shown. Each component of the application is represented by an individual icon. The lower-left window shows the contents of one of these components— an instantiation of the ERO created using Ptolemy's synchronous dataflow domain. The central window shows the contents of another application component— the source code defining an ICP action object for performing exponential rate estimation by using the ERO on a single-tone component selected from the system blackboard. A graphical blackboard browser, created using the Tcl/Tk language embedded within Ptolemy, is displayed in the rightmost window. This browser allows interactive inspection of the system's internal state at run-time.

7 CONCLUSION

We have described a complex multi-tone transient detection problem and its solution based on utilizing symbolic heuristic search along with numerical signal processing. This solution serves to highlight the need for signal processing design environments which support the systematic integration of symbolic heuristic search with numerical signal processing. In this paper, we have described one such environment which has been recently developed by incorporating the ICP tool for blackboard-based heuristic search within the heterogeneous Ptolemy environment. This environment, which is known as the IPUS domain for Ptolemy, is currently being used for implementing our solution to the multi-tone transient detection problem.

8 ACKNOWLEDGMENTS

This work was sponsored in part by the Department of the Navy, Office of the Chief of Naval Research, contract number N00014-93-1-0686 as part of the Advanced Research Projects Agency's RASSP program. The authors wish to acknowledge Kuang-Jung Huang and Kwong Kit Wong for their contributions to the implementation of the IPUS domain in Ptolemy.

9 REFERENCES

[1] N. Carver and V. Lesser, "Blackboard Systems for Knowledge-Based Signal Understanding," pp. 205-250, in *Symbolic and Knowledge-Based Signal Processing*, edited by A.V. Oppenheim and S.H. Nawab, Prentice Hall, Englewood Cliffs, N.J., 1992.

[2] J. M. Winograd and S. H. Nawab, "A C++ Software Environment for the Development of Embedded Signal Processing Systems," *Proc. ICASSP 1995*, vol. 4, pp. 2715-19, Detroit, May 1995.

[3] J. M. Winograd, "IPUS C++ Platform User's Manual," *Tech. Report, Rome Laboratory Air Force Material Command*, Griffiss Air Force Base, NY, 1994.

[4] J. T. Buck, S. Ha, E. A. Lee, and D. G. Messerschmitt, "Ptolemy: A Framework for Simulating and Prototyping Heterogeneous Systems," *Int. Journal of Computer Simulation*, special issue on "Simulation Software Development," vol. 4, pp. 155-182, April, 1994.

[5] R. Mani and S. H. Nawab, "A Multiband Exponential Rate Operator for Musical Transient Analysis," *Proc. ICASSP 1995*, vol. 2, pp. 1233-36, Detroit, May, 1995.

[6] H. M. Hanson, P. Maragos, and A. Potamianos, "A System for Finding Speech Formants and Modulations Via Energy Separation," *IEEE Trans. Speech Aud. Proc.*, vol. 2, no. 3, July, 1994.

[7] P. Maragos, T. F. Quatieri, and J. F. Kaiser, "On Amplitude and Frequency Demodulation Using Energy Operators," *IEEE Trans. Sig. Proc.*, vol. 41, pp. 1532-50, Apr., 1993.

[8] D. L. Jones and T. W. Parks, "A High Resolution Data-Adaptive Time-Frequency Representation," *IEEE Trans. Acoust. Speech Sig. Proc.*, vol. 38, no. 12, pp. 2127-35, Dec. 1990.

[9] V. R. Lesser, S. H. Nawab, and F. I. Klassner, "IPUS: An Architecture for the Integrated Processing and Understanding of Signals," *Artificial Intelligence*, vol. 77, pp 129-171, 1995.

[10] The MathWorks, Inc., Natick, MA, *Matlab User's Guide*, Aug. 1992.

[11] E. E. Milios and S. H. Nawab, "Signal abstractions in signal processing software," *IEEE Trans. Acoust., Speech, and Signal Processing*, vol. 37, pp. 913–928, June 1989.

[12] N. Carver and V. Lesser, "A planner for the control of problem-solving systems," *IEEE Trans. Sys., Man, and Cybernetics*, vol. 23, no. 6, pp. 1519–1536, 1993.