

Never Reprogram Again™

Ted J. Biggerstaff
Software Generators, LLC
Austin, Texas USA
dslgen at softwaregenerators dot com

Abstract—DSLGen™ (Domain Specific Language Generator) is a program generation system in which application programs can be written in a domain specific language that is independent of the execution platform architecture and yet can be targeted to arbitrary existing and future execution platforms in a way that exploits the performance or computation improvement opportunities specific to those platforms. This allows switching from one execution platform to another without reprogramming the applications. The generation of target programs is fully automatic and requires no user input or action beyond the specification of the computation and the separate specification of the features of the target execution platform.

Keywords -- *associative programming constraints, natural and synthetic partitions, design patterns, logical and physical architectures, design feature encapsulation, implementation neutral specification, domain specific languages, inference, problem domain inference, partial evaluation.*

I. INTRODUCTION

DSLGen™ (patents issued [8] [9] [10]) is a transformation-based program generation system that fully automatically generates a target implementation from two independent specifications: 1) a domain specific, *Implementation Neutral Specification (INS)* of the desired computation and 2) a domain specific *EXecution Platform Specification (EXPS)* that describes the features of the execution platform upon which the application code will run. **The INS is invariant over target execution platform architectures.** That is, an application programmer can make no predictions about the architecture of the target implementation by looking at the INS alone. Thus, no reprogramming of the INS is required to switch from one platform to another. Only the EXPS features need to be changed to switch from one architecture (e.g., multicore) to another (e.g., vector machines). Importantly, DSLGen™ fully automatically converts an INS and an EXPS into target implementation code that takes advantage of a broad range of opportunities for high capability computations including large grain parallelism (e.g., multicore CPUs), small grain parallelism (e.g., instruction level parallelism or ILP), design pattern frameworks and so forth. It is theoretically possible to extend DSLGen™'s capabilities to other target execution platforms such as GPUs, Digital Signal Processors (DSPs), specialized processors, Field Programmable Gate Arrays (FPGAs), and API interfaces to layered implementations or libraries. The author believes that DSLGen™ can be extended with new transform sets that will produce output optimized for virtually any arbitrary existing or future

architecture. How can DSLGen™ automatically produce programs that are tailored to such highly varied execution architectures?

The short answer is that DSLGen™ is an extensible generator that is designed to create a program design from scratch based on the INS plus generalized constraints and design features specified in the EXPS. In some sense, it is doing what a human programmer does. DSLGen™ automatically builds a Logical Architecture (LA) that constrains some problem domain oriented features of the target program design but defers building a Physical Architecture (PA) that commits to programming language oriented features (e.g., routine architectures, parametric connections, communication patterns and synchronization patterns). That is, DSLGen™ architects, designs, constrains, reorganizes and optimizes the target program in the problem and programming process domains rather than in the programming language (PL) domain and only after the macroscopic structure of the program is settled does it generate PL code. In short, it designs the solution first and codes it second.

Part of the secret to this process is that DSLGen™ eschews PL representations during the design and architecture portion of the process thereby freeing it from the highly restrictive constraints of PLs. PLs are solution oriented not design or architecture oriented. They require the programmer to tell how to do a computation whereas during these early phases, the programmer knows what needs to be done and what design features the solution will have (i.e., the computational goals) but has not yet fully determined how to implement and integrate the computational needs and solution features.

II. THE PROBLEM

A key problem in exploiting the capabilities of various existing and future execution platform architectures for a specific target computation is the conflict between the goal of precisely describing the implementation of a target computation and the goal of casting the implementation into a variety of forms each of which exploits a different set of high capability features of some specific execution platform architecture (e.g., parallel processing via multicore based threads). The key culprit in this conflict is the representation system used in the course of creating a target program – that is, the use of programming language based abstractions to represent the evolving program at each stage of its development. Einstein said “We see what our languages allow us to see.” And when a computer scientist understands his or her world in terms of programming languages, it is natural to construct intermediate design and precursor

representations in terms of programming language based abstractions. This has led to our conventional, reductionist, top-down models of program design and development, which the author believes has been a key impediment to mapping an implementation neutral specification of a computation to an arbitrary platform while still exploiting whatever high capability features that platform possesses.

In such a top-down model, the structure and some details of a layer of the target program are specified along with some abstract representation of the constituent elements (i.e., lower level layers) of that layer. In human based application of top-down design, the abstract elements of the lower level layers are often expressed in terms of an informal pseudo-code. In the automated versions of top-down design, the pseudo-code is often replaced by formal expressions (i.e., programming language based expressions) of the interfaces to the lower level layers, which may be simple PL calls, object oriented invocations, or skeletal forms of elements that remain to be defined. Alternatively, these interfaces may be calls or invocations to fully defined API layers or interfaces to message based protocols (e.g., finite state machine specifications). In any case, the structure is fixed at a high level before the implications of that structure become manifest in a lower level, later in the development process. Refinements within the lower layers often require changing or revising the structure at a higher level, which can be problematic. Further, in an automated system, distinct programming design goals will be, by necessity, handled at different times. This is further complicated by the fact that multiple design goals may be inconsistent (at some level of detail) or at least, they may be difficult to harmonize.

A good example of this kind of difficulty is trying to design a program to exploit thread based parallel implementation. The exact structure and details of the final program are subtly affected by a myriad of possible problem features and programming goals. A threaded implementation will require some thread synchronization logic which may be spread across a number of yet to be defined routines. The computation will have to be partitioned into parts that are largely determined by the specifics of the target computation. These partitions will be mapped into routines and threads (e.g., some lightweight computations batched in one thread and other heavyweight computations decomposed into slices with their own threads). The thread protocol will introduce low level implementation details that potentially will have to be harmonized across a number of routines. The parameter choices for these routines (i.e., the plumbing) may be involved in the communication design for these thread routines and will be constrained by low level implementation details of the thread protocol. In DSLGenTM, such programming language level routine structures, routine inter-communication decisions, thread protocol restrictions and thread library implementation requirements are added into the architecture close to the end of the design process.

If an automated generator tries to handle all of these design issues at once, there is an overwhelming explosion of cases to deal with and the approach quickly becomes infeasible.

III. THE SOLUTION

The ideal solution would be to recognize design goals and assert the programming process objectives provisionally (e.g., organize the computation to exploit threads) without committing fully and early-on to constructing the PL structures and details, because those PL structures and details are likely to change and evolve as the target program is refined toward a final implementation. The ideal solution would allow each design issue or feature to be handled atomically, one at a time. Then, if necessary, those previously asserted provisional commitments could be altered before they are cast into concrete code. And this is the essence of DSLGenTM.

DSLGenTM allows the construction of a logical architecture that levies minimal constraints on the evolving program and explicitly defers generating programming language expressions early on. That is, initially the LA will constrain only the decomposition of a computation into its major (and natural) organizational divisions (which are called *natural partitions*) omitting any PL details of the programming routine structure or PL details of those major organizational divisions. There is no information on control structure, routines, functions, threads, parametric connections, data flow connections, machine units, instruction styles, parallel synchronization structures and so forth. All of that is deferred and added step by step as the generation process proceeds. In fact, the LA will be revised and evolved step by step by the encapsulation of individual design features, each of which will further constrain the final expression of the target program.

A. Associative Programming Constraints and the LA

DSLGenTM builds the LA out of a new kind of representation element – an *Associative Programming Constraint (APC)*. APCs are partial and provisional constraints on the target computation. They do not fully determine the target implementation. These APCs come in two major varieties: *Iteration constraints* and *partition constraints*. For example, a *loop constraint* (a subclass of iteration constraint) might specify “i” and “j” to be indexes of a matrix “a” that have ranges of [0,(m-1)] and [0,(n-1)], respectively. And related to this loop constraint, for example, might be a partition constraint (e.g., Edge1) that specifies the subdivision of that loop constraint for which (i==0). Nothing more about the implementation is determined by these constraints.

Operationally, APCs are CommonLisp Object System (CLOS) objects that are associated with elements of the INS and initially arise via translation of the INS. They are logical in the sense that their essential specification mechanism is based on predicate logic assertions. These assertions will be altered and extended as the generation process proceeds thereby altering and refining the definitions of the constraints. APCs are propagated over the INS structure (somewhat analogous to APL’s method of loop introduction and placement). They are combined in several ways, the operational effect of which is to merge equivalent iterations or to adapt two slightly different computational cases to a single iteration scheme. They can be split in two,

reorganized into groups that imply future design features and revised to incorporate one or more design features. Not until later are they actually applied, decomposing the INS into specialized subdivisions of the implementation, which creates the precursors to actual code.

Specialize versions of these two major classes of APCs may be created by subclassing, thereby allowing other kinds of architectural factorings. To provide a concrete context in which to discuss the LA and its representational elements, we will introduce a problem domain and a domain specific language for that problem domain.

IV. THE PROBLEM DOMAIN AND AN EXAMPLE PROBLEM

The initial problem domain treated by DSLGen™ is digital signal processing (DSP) and includes problems that range from signal and image processing to neural networks to pattern recognition plus a rich set of related problems. The domain specific language used to express the INS is based on the Image Algebra (IA) [27].

As an example computation, we develop a program that performs Sobel edge detection on a grayscale image (i.e., where the pixels are shades of gray). Such a program would take, for example, the image “a” in Fig. 1 as input and produce the image “b” in Fig.2 as output. The output image has been processed so as to enhance (line) edges of items in the image by the Sobel edge detection method.

Each black and white pixel $b[i,j]$ in the output image “b” is computed from an expression involving the sum of products of pixels in a neighborhood (e.g., sp , of type *iemplate*) surrounding the $a[i,j]$ pixel and the coefficients defined by that neighborhood (e.g., s). This is called a *convolution* of a matrix with a template (or neighborhood). In the IA, a convolution is designated by the \oplus operator, e.g., $(a \oplus s)$. In the following examples, s and sp will designate instances of the class *iemplate*. Mathematically, the Sobel computation is defined as

$$\{\text{Forall}_{i,j} \{b_{i,j} : b_{i,j} = \sqrt{(\sum_{p,q} (w(s)_{p,q} * a_{i+p,j+q})^2 + \sum_{p,q} (w(sp)_{p,q} * a_{i+p,j+q})^2)}\} \quad (1)$$

where i and j are indexes that range over the matrices a and b ; p and q are indexes that range over the *iemplate* neighborhoods s and sp ; and the coefficients of the neighborhood (which are also called *weights*) are defined by the function “ w ”. For Sobel edge detection, the weights are all defined to be 0 if the center pixel of the neighborhood corresponds to an edge pixel in the image (i.e., $w(s) = 0$ and $w(sp) = 0$), and if not an edge pixel, they are defined by the s and sp neighborhoods shown in (2). It is convenient to index the neighborhoods in the DSL from -1 to +1 for both dimensions so that the current pixel being processed is at (0, 0) of the neighborhood.

$$w(s) = P \begin{matrix} & \overbrace{\begin{matrix} -1 & 0 & 1 \end{matrix}}^Q \\ \begin{matrix} -1 \\ 0 \\ 1 \end{matrix} \left[\begin{matrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{matrix} \right] \end{matrix} \quad w(sp) = P \begin{matrix} & \overbrace{\begin{matrix} -1 & 0 & 1 \end{matrix}}^Q \\ \begin{matrix} -1 \\ 0 \\ 1 \end{matrix} \left[\begin{matrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{matrix} \right] \end{matrix} \quad (2)$$

Since an implementation of this computation for a parallel computer may not be organized like the mathematical formula, it is useful to represent this specification more abstractly because such abstractions can defer the implementation and organization decisions and thereby allow the computation (i.e., what is to be computed) to be specified completely separately and somewhat independently from the implementation form (i.e., how it is to be computed). Thus, the abstract computation specification is independent of the architecture of the machine that will eventually be chosen to run the code. Choosing a different machine architecture for the implementation form without making any changes to the specification of the computation (i.e., the what), will automatically generate a different implementation form that is tailored to the new machine’s architecture. More to the point, porting from one kind of machine architecture (e.g., machines with instruction level parallelism like Intel’s SSE instructions) to a different kind of machine architecture (e.g., machines with large grain parallelism such as multi-core CPUs) can be done automatically by only making trivial changes to the machine specifications and no changes to the computation specification (i.e., the what). The publication form in [27] for the Sobel Edge detection mathematical formula (1) is based on the Image Algebra domain specific language (DSL). Re-expressing the formula (1) in the Image Algebra gives a first cut at the INS for the Sobel example:

$$b = [(a \oplus s)^2 + (a \oplus sp)^2]^{1/2} \quad (3a)$$

Of course, the INS will need some declarations for a , b , s , sp , etc.:

```
(DSDeclare IATemplate s :form (array (-1 1) (-1 1))
      :of DSNumber)
(DSDeclare IATemplate sp :form (array (-1 1) (-1 1))
      :of DSNumber)
(DSDeclare DSNumber m :facts ((> m 1)))
(DSDeclare DSNumber n :facts ((> n 1)))
(DSDeclare BWImage a :form (array m n) :of BWPixel)
(DSDeclare BWImage b :form (array m n) :of BWPixel)
b = [(a \oplus s)^2 + (a \oplus sp)^2]^{1/2} \quad (3b)
```

m and n are assumed to be user defined. The DSL type declarations (e.g., *IATemplate*, *BWImage*, etc.) define CLOS types that will eventually refine to C types. The “:facts” keyword denotes a conjunction (i.e., list) of facts pertinent to the declared item (e.g., m) and will be used to infer, for example, that“($i==(m-1)$)” is false when “($i==0$)” is true. Beyond (3b), we will also need some definitions for s and sp equivalent to (2) and for \oplus , all to be defined later.

This DSL is the basis of the Implementation Neutral Specification (INS) in the examples used throughout the remainder of this document. A full description of the IA used by DSLGen™ is beyond the scope of this paper (see [27]) but a few comments are in order. The IA is much like APL in the sense that IA specifications eschew the use of explicit looping constructs allowing loops to be implied by IA

operators and data structures. The generator will introduce implied loops as constraints and, through the manipulation, combination and propagation of these constraints, will determine the relationships between IA expressions and loops. The initial form of the LA arises during this process.



Figure 1. Input Image a

In DSLGenTM, the Image Algebra is adapted to a more utilitarian, LISP based syntax with prefix operators, without the pretty symbols (e.g., the convolution operator \oplus becomes a Lisp symbol), and with the w functions in (1) becoming so-called *Method-Transforms (MT)*, which rewrite *Abstract Syntax Tree (AST)* subtrees. MTs look superficially a bit like object oriented methods with a pattern (i.e., the MT's *left hand side* or *lhs*) as the analog of a method's parameter sequence and a pure functional expression *right hand side (rhs)* as the analog of a method's body. MTs will be an important component of the *intermediate language (IL)* by which provisional but malleable definitions are expressed. For example, w of the neighborhood s is an MT expressed as:

```
(Defcomponent w (sp #. ArrayReference ?p ?q)
  (if (or (== ?i ?ilow) (== ?j ?jlow)
        (== ?i ?ihigh) (== ?j ?jhigh)
        (tags (constraints partitionmatrixtest edge)))
    (then 0)
    (else (if (and (!= ?p 0) (!= ?q 0))
              (then ?q)
              (else (if (and (== ?p 0) (!= ?q 0))
                        (then (* 2 ?q))
                        (else 0)))))))) (4)
```

where *ArrayReference* is the name of a shared pattern that will recognize an array reference in an AST (e.g., $a[i,j]$) and bind the loop index variables (e.g., i and j) to the pattern variables $?i$ and $?j$, the matrix name a to $?a$ and the expressions defining the upper and lower ranges of those loop indexes to $?ihigh$, $?ilow$, etc. The remainder of the *lhs* pattern after *ArrayReference* will bind $?p$ and $?q$ to the loop index names used by the inner convolution loops over the neighborhood designated by sp .

The “tags” expression designates a property list for the OR conditional expression, which in (4) provides the user supplied domain knowledge that the OR expression is a *partitioning condition* for this computation that will identify *edge* partitions and by implication, a non-edge (i.e., *center*) partition. Problem domain concepts like “edge” and “center” play a key role in the logical architecture for the target computation and beyond that, in imposing design pattern frameworks onto a logical architecture. Heuristic rules based on domain concepts are the mechanisms whereby DSLGenTM chooses a design pattern framework to introduce PL structures and clichés (e.g., coordinated routines, synchronization patterns and thread management clichés) and maps the LA into the structures and clichés of that design pattern framework.




Figure 2. Output Image b

The opportunity for such domain specific heuristic rules is open ended, especially given the rich variety of possible semantic subclasses of partitions. Different problem examples may introduce other domain semantics. For example, in the matrix domain, the semantic subclasses include *corners* (e.g., corners are special cases in partitioning image averaging computations); *non-corner edges* also used in image averaging; *upper and lower triangular matrices*, which are used in various matrix algorithms; *diagonal matrices*; and so forth. By contrast, in the data structure domain, domain subclasses include *trees*, *left and right subtrees*, *red and black nodes*, etc. In general, domain concepts drive the DSLGenTM program generation process.

V. THE DESIGN REPRESENTATION SYSTEM

The first iteration of the logical architecture for the Sobel example is shown conceptually in Figure 3. Loop constraints are CLOS objects that keep track of loop indexes, loop nesting and the logical description of the loop, which comprises logical assertions and precursors thereof. For example, *Partestx* of s is IL manufactured during INS reduction. It is a precursor to a logical assertion that will refine to a partitioning condition for some (not yet decided upon) partition. *Partestx* of s will eventually be refined to a concrete expression such as “ $(i==0)$ ” in the context of a

particular partition-based computation (e.g., Edge1). And the addition of “(i==0)” to the loop constraint will change the form of the C code that is eventually generated for that partition by causing the loop over “i” to evaporate and possibly allowing the body of the loop to be simplified. In the chosen example, the bodies of edge loops undergo significant simplification.

Operationally, Partestx is a closure over one of the disjuncts (e.g., (i==0)) in the OR expression in (4) and the translation context bindings (e.g., ((?i i) (?ilow 0)) at the time of Partestx formation. That translation time will be when an expression like “(a ⊕ s)” is being translated and a provisional loop constraint is being introduced and propagated to the “⊕” level expression. As loop constraints are introduced, propagated and combined (e.g., providing loop sharing for separate computations), DSLGen™ provides machinery for recording design decisions (e.g., discarding unneeded loop indexes) via dynamically generated transformations that will be applied periodically to synchronize the overall design.

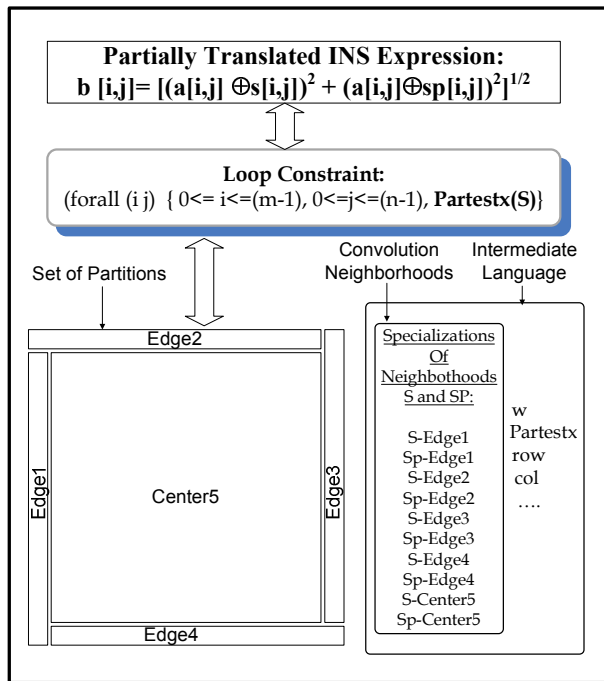


Figure 3. Initial logical architecture of example

The loop constraint is associated with a partially translated INS expression (by appearing on the INS’s tags list). Generally speaking, the loop constraint may be associated with a set of partitioning constraints such as the Edge1, Edge2, Edge3, Edge4 and Center5 (i.e., the CLOS objects) of this example. They indicate a partial and provisional decomposition of the loop, where each decomposition body eventually will be formed from a cloned and specialized version of the associated INS expression. But DSLGen™ does not perform the decomposition yet, because as the implementation design evolves, the partitioning is almost certain to change before it is cast into code. The

partitioning implied by the set of partition objects is sort of a “to do” list and a “to do” list that will likely change before it is turned into code. However, this future cloning and specialization will be accomplished by using a set of newly formed specializations of s, sp and their IL. For example, the specialization of a specific neighborhood (e.g., sp) and its IL (e.g., w) for a specific partition constraint (e.g., edge1) is formed by assuming a truth value for the partitioning condition of the partition constraint and partially evaluating the IL definitions under that assumption. For example, for Edge1, the MT definition of w of sp, in (4), would partially evaluate to a new MT definition, w of sp-edge1:

$$(\text{Defcomponent } w \text{ (sp-edge1 \#.ArrayReference ?p ?q) 0) (5)$$

The LA is malleable so that DSLGen™ can incrementally introduce design features by a process called *Design Feature Encapsulation (DFE)*. DFE will revise IL definitions, extend and reorganize partition sets and occasionally even revise some of the DSLGen™’s own transformations that define the overall generation and programming process (e.g., when introducing instruction level parallelism).

A. Design Feature Encapsulation

For our example, let us use an EXPS of “((PL C) Mcore (Threads MS) (LoadLevel (SliceSize 5)))” where C is the output language, the target is a multicore machine that exploits threaded parallelism using Microsoft’s thread library and the design should decompose the computation by slicing up some unspecified heavyweight computation using 5 unspecified units per slice. In the example, the LA specifics will be used to disambiguate what is being sliced up (e.g., Center5) and what the units are (e.g., matrix rows).

In figure 3, we have already seen a simple example of DFE where IL definitions are specialized to specific logical partitions of a target computation. These specializations will cause computations along the matrix edges to simplify to a single loop that assigns 0 to pixels of that edge. Another simple example of DFE is mapping from IA neighborhood style indexing to C style indexing. IA style indexing ranges from -n to +n for an (2n+1) by (2n+1) neighborhood so that the center pixel is at (0,0). In contrast, the C language (i.e., the chosen output language) ranges from 0 to 2n. The indexing DFE is accomplished by algebraic manipulation of the right hand side (i.e., the MT body) of IL involving neighborhood loop indexes, which relocates instances of those loop indexes appropriately.

However, one of the most powerful examples of DFE is the introduction of architectural design features that alter the form of and relationships within the implementation across a broad set of coordinated routines, data structures and possibly even parallel processes. This is accomplished by the use of *synthetic partitions*, which extend the notion of natural partitions by adding implied design feature constraints.

1) Synthetic Partitions

In DSLGen™, the generation process is divided into named phases, each of which has a narrowly defined

generation purpose. The phase most relevant to the introduction of wide ranging design features is the Synthetic-Partitioning phase. During the SyntheticPartitioning phase, the generator introduces design features (via synthetic partition objects) that will constrain the evolving LA to be much more specific to a design for some execution platform. These synthetic partitions imply implementation structures that exploit high capability features of the execution platform and that, when finally re-expressed in a form closer to code, may have wide ranging and coordinated affect across much of the LA (e.g., via multiple routines that coordinate the use of multicore parallel computation). The SyntheticPartitioning phase operates on the logical architecture to reorganize the partitions and probably (depending on the execution platform spec) create synthetic partitions that connect to one or more code frameworks. These code frameworks hold the implementation details (e.g., thread and synchronization management) to be integrated into the evolving target program. The synthesis process for this example includes the following detailed steps.

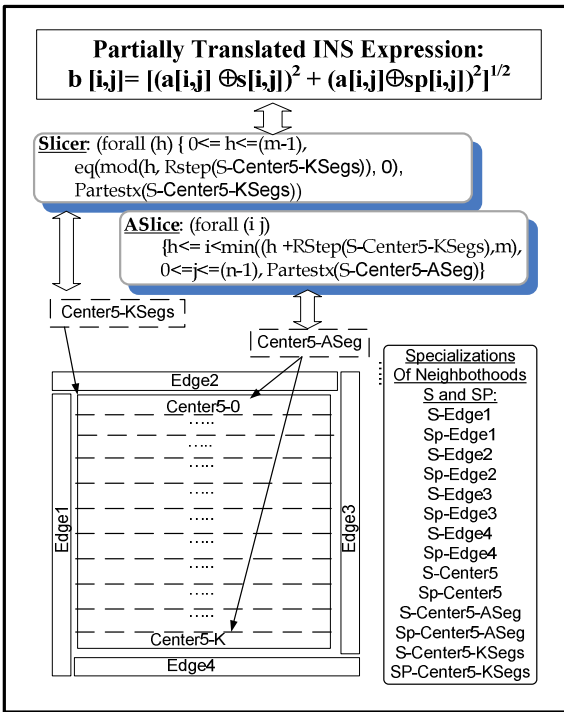


Figure 4. Revised logical architecture of example

Let us say that the EXPS requires that the computation should be load leveled (i.e., sliced into smaller computational pieces) in anticipation of formulating the computation to run in parallel threads on a multicore platform, which we have also included in the EXPS. Given these assumptions, Fig. 4 shows the revised logical architecture for these assumptions (with synthetic partitions denoted by dashed boxes). Load leveling will introduce two synthetic partitions (e.g., Center5-KSegs and Center5-ASeg) that respectively express the design feature that decomposes the center partition (i.e.,

Center5) into smaller pieces and the design feature that processes each of those smaller pieces.

Simultaneously, in Fig. 4, the loop constraint from Fig. 3, is reformulated into two loop constraints (i.e., Slicer and ASlice) that will be required by the synthetic partitions Center5-KSegs and Center5-ASeg. This synthesis process also introduces versions of the neighborhoods S-Center5 and SP-Center5 specialized for Center5-KSegs and Center5-ASeg and generates specialized IL for each. The step size of the Slicer loop is inferred from information in the EXPS or by a default if the EXPS is silent on the subject. It is represented by the IL expression “Rstep(S-Center5-KSegs)” in Fig. 4. For the example, we have chosen a step size of 5. Using this step size, Slicer will dynamically compute a new range for each instance of the ASlice loop.

2) Imposing Design Patterns on a Logical Architecture

Now, DSLGen™ is ready to add in the PL level details (e.g., sets of interrelated routines, parametric plumbing, thread management clichés and protocols of specific thread libraries) by mapping the LA into a PA through use of a design pattern framework. To make the example interesting, we have assumed a design with thread-based multicore parallelism (in the EXPS).

DSLGen™ allows for a library of design pattern based frameworks (i.e., objects with associated PL-like skeletons), each of which represents some reasonably small combination of related design features. Additionally, each such framework has a set of holes (indicated by **bolded designators**) that are tailored to the LA’s combination of architectural features. These holes are designed to receive computational payloads from the LA (e.g., partition specific computations). For example, a particular framework might be designed to receive partitions such as image edges that are “probably” order n computations (i.e., lightweight computations) as well as to receive partitions such as image centers that are “probably” order n squared computations (i.e., heavyweight computations). Such a framework might introduce a set of cooperating PL routines and the parametric plumbing among those routines, where the plumbing may include some “holes” that will receive data items specific to the INS. There may be additional PL design features included, such as synchronization patterns for parallel computation and detailed thread control clichés. But the framework is agnostic about its payload. It says nothing about exactly what kind of a computation is occurring in its holes. That computational payload information will be supplied by the logical architecture.

So, based on the example LA plus specific features required by the EXPS, DSLGen™ will search its design pattern data base for a design pattern meeting these criteria. It finds one with the following skeletal PL framework:

```
void ?managethreads ( )
{ HANDLE threadPtrs[200];
  HANDLE handle;
  /* Launch the thread for lightweight processes. */
  handle = (HANDLE)_beginthread(
    &?DoOrderNCases , 0, (void*) 0);
```

```

DuplicateHandle(GetCurrentProcess(), handle,
    GetCurrentProcess(),&threadPtrs[0],
    0, FALSE, DUPLICATE_SAME_ACCESS);
/* Launch the threads for the slices of heavyweight
processes. */
{handle = (HANDLE)_beginthread(& ?DoASlice , 0,
    (int) (Index ?SlicerConstraint) ) ;
DuplicateHandle(GetCurrentProcess(), handle,
    GetCurrentProcess(),&threadPtrs[tc],
    0, FALSE, DUPLICATE_SAME_ACCESS);
tc++; } (tags (constraints ?SlicerConstraint))
long result = WaitForMultipleObjects(tc, threadPtrs,
    true, INFINITE); } (6)

```

```

void ?DoASlice (int (Index ?SlicerConstraint))
    { { ?ins } (tags (constraints ?ASliceConstraint))
      _endthread(); } (7)

void ?DoOrderNCases ( )
    { ?OrderNCases
      _endthread(); } (8)

```

Associated with the class of this design pattern is a CLOS method whose job is to find key elements in the LA and bind them to pattern variables (e.g., **?ins** and **?SlicerConstraint**); invent and bind unique names for routines (e.g., “SobelCenter8” might be invented for **?managethreads**); clone and specialize the INS to specific partitions (e.g., by substituting sp-Edge1 for sp); and instantiate the skeletons with the bindings. Notice that the design skeletons are agnostic as to what their computational payload is going to be. Further, there are no PL like connections (e.g., calls to PL routines) between the design pattern skeletons and anything in the LA. The only requirements of the design pattern are that the LA has partitions that represent lightweight processes that can be batched in a single thread (e.g., edges) and a heavyweight process (e.g., a center) that is partitioned into a slicer partition and an implied set of slice partitions. These requirements are determined by **domain logic**, that is, logical rules operating on problem domain information (e.g., properties of edges) rather than PL information.

Space limitations preclude showing the full step by step expansion of all these skeletal routines but the thread routine that batches the edge partitions (**?DoOrderNCases**) is reasonably short and is interesting in that the edge loops will drastically simplify when in-lined and partially evaluated. Instantiating with cloning and specialization produces:

```

void SobelEdges9( )
{ /* Edge1 partitioning condition is (i=0) */
  for (int j=0; j<=(n-1);++j)
    b [0,j]= [(a[0,j] ⊕ s-edge1[0,j])2 +
              (a[0,j] ⊕ sp-edge1[0,j])2]1/2 }
/* Edge2 partitioning condition is (j=0) */
for (int i=0; i<=(m-1);++i)
  b [i,0]= [(a[i,0] ⊕ s-edge2[i,0])2 +
            (a[i,0] ⊕ sp-edge2[i,0])2]1/2 }

```

```

/* Edge3 partitioning condition is (i=(m-1)) */
for (int j=0; j<=(n-1);++j)
  b [(m-1),j]= [(a[(m-1),j] ⊕ s-edge3[(m-1),j])2 +
                (a[(m-1),j] ⊕ sp-edge3[(m-1),j])2]1/2 }
/* Edge4 partitioning condition is (i=(n-1)) */
for (int i=0; i<=(m-1);++i)
  b [i, (n-1)]= [(a[i, (n-1)] ⊕ s-edge4[i, (n-1)])2 +
                 (a[i, (n-1)] ⊕ sp-edge4[i, (n-1)])2]1/2 }
_endthread(); } (9)

```

Notice that in (9) partial evaluation plus inference has caused one of each pair of the edge loops in (9) to evaporate and zeros appear in one of the index positions in the array expressions. In truth, these loop refinements occur concurrently with the inlining of the IL definitions (see the following section) but in the name of space, showing it here shortens (9) and makes them easier for the reader to understand.

While the expansion of the **?DoASlice** routine is longer than **SobelEdges9**, it is important because it shows the default partition specialization (i.e., the center slice partition). It is populated with the **ASlice** loop constraint plus the INS specialized to **S-Center5-ASeg** and **SP-Center5-ASeg**. The example code is shown with a slice size of 5 but alternatively, it could be declared by the user to be a parameter. Before inlining the IL definitions, the **?DoASlice** routine is specialized to:

```

void SobelCenterSlice10 (int h)
{
  for (int i=h; i<= min((h+ 4),(m-1)); ++i)
    {for (int j=1; j<=(n-2); ++j)
      {b [i,j]= [(a[i,j] ⊕ s-center5-ASeg[i,j])2 +
                (a[i,j] ⊕ sp-center5-ASeg[i,j])2]1/2 } }
    _endthread(); } (10)

```

Like (9), form (10) shows the loop refinements out of order to save space and make (10) easier to understand. The range of j is now [1,(n-2)] rather than [0,(n-1)] because of the effect of the partitioning condition.

3) Inlining Intermediate Language Definitions

The DSLGenTM Inlining phase will inline the IL definitions, replacing the convolution expressions with their definitions such as that of the convolution operator, i.e.,

```

(* (a[(row sp-Edge1 a[i,j] p q), (col sp-Edge1 a[i,j] p q)])
  (w sp-Edge1 a[i,j] p q))2 (11)

```

for each partition specific INS clone. The inlining will continue recursively for the lower level IL definitions, e.g., row, col and w (where row and col map from neighborhood coordinates to matrix coordinates). Since (w sp-Edge1 a[i,j] p q) is defined as 0 in (5), expression (11) partially evaluates to 0. Similarly, all other convolution expressions involving edges partially evaluate to 0. After all of the inlining and partial evaluation (but before adding local declarations), expression (9) becomes (12):

```

void Sobel Edges9( )
{ /* Edge1 partitioning condition is (i=0) */
  {for (int j=0; j<=(n-1);++j) b [0,j]=0;}
  /* Edge2 partitioning condition is (j=0) */
  {for (int i=0; i<=(m-1);++i) b [i,0]= b [0,j]=0;}
  /* Edge3 partitioning condition is (i=(m-1)) */
  {for (int j=0; j<=(n-1);++j) b [(m-1),j]=0;}
  /* Edge4 partitioning condition is (i=(n-1)) */
  {for (int i=0; i<=(m-1);++i) b [i, (n-1)]= 0;}
  _endthread( ); } (12)

```

And analogously for the **?DoASlice** routine, after a series of inlining steps analogous to the Edge1 partition refinement process but without the extensive simplification engendered by the IL definitions for the edge partitions, the center slice partition case (10) refines into:

```

void SobelCenterSlice10 (int h)
{long ANS45; long ANS46;
  /* Center5-KSegs partitioning condition is
    (and (not (i=0)) (not (j=0)) (not (i=(m-1)))
      (not (j=(n-1)))) */
  /* Center5-ASeg partitioning condition is
    (and (not (i=0)) (not (j=0)) (not (i=(m-1)))
      (not (j=(n-1))) (h<=i) (i<=(min (h+4) (m-1))))*/
  for (int i=h; i<=min((h+ 4),(m-1)); ++i) {
    for (int j=1; j<=(n-2); ++j) {
      ANS45 = 0;
      ANS46 = 0;
      for (int p=0; p<=2; ++p) {
        for (int q=0; q<=2; ++q) {
          ANS45 +=
            (((*(b + ((i + (p + -1)))))) + (j + (q + -1)))) *
            (((p - 1) != 0) && ((q - 1) != 0)) ? (p - 1):
            (((p - 1) != 0) && ((q - 1) == 0)) ?
              (2 * (p - 1)): 0));
          ANS46 +=
            (((*(b + ((i + (p + -1)))))) + (j + (q + -1)))) *
            (((p - 1) != 0) && ((q - 1) != 0)) ? (q - 1):
            (((p - 1) == 0) && ((q - 1) != 0)) ?
              (2 * (q - 1)): 0)); } }
    int i1 = ISQRT ((pow ((ANS46), 2) +
      pow ((ANS45), 2)));
    i1 = (i1 < 0) ? 0 : ((i1 > 0xFFFF) ? 0xFFFF : i1);
    ((*((A + (i))) + j)) = (BWPIXEL) i1; } }
  _endthread( ); } (13)

```

The examples are adapted from generated code to accommodate the format and space available. For example, in generated code, i and j would be generated names like idx3 and idx4. Similarly, p and q would be something like p15 and q16. Additionally, in (13), a discussion of the introduction of the answer variables (e.g., ANS45) and the masking expression near the end is beyond the scope of this paper.

The reader will note that the inlining step has introduced some common sub-expressions (e.g., (p - 1)) which will degrade the overall performance if not removed. If this code

is targeted to a good optimizing compiler, these common sub-expressions will be removed by that compiler and thereby the performance improved. However, if the target compiler is not able to perform this task, DSLGen™ offers the option of having the generator system remove the common sub-expressions and this can be easily added to the specification of the execution platform. However, the common sub-expressions are explicitly included in this example (i.e., not optimized away) to make the connection to the structures of the MTs used by the INS more obvious to the reader. The broad structure of the right hand operand of the times (*) operator in the right hand side of the assignments to the answer variables ANS45 and ANS46 is structurally the same as that of the W method transform specialized to the center partition for SP and S. That is, the right hand side of the C form:

```

((((*(b + ((i + (p + -1)))))) + (j + (q + -1)))) *
  (((p - 1) != 0) && ((q - 1) != 0)) ? (q - 1):
  (((p - 1) == 0) && ((q - 1) != 0)) ? (2 * (q - 1)): 0)) (14)

```

mimics the form of the MT definition for w of SP-Center5 because (14) is derived by inlining that MT definition and eventually processing it into legal C. For reference, the rhs of the MT definition of w of SP-Center5 has the form

```

(if (and (!= ?p 0) (!= ?q 0))
  (then ?q)
  (else (if (and (== ?p 0) (!= ?q 0))
    (then (* 2 ?q))
    (else 0))))). (15)

```

When the inlining occurs, the SP-Center5 generator pattern variable ?p is bound to “(p - 1)” and ?q is bound to “(q - 1)”. The “- 1” part of these values arise because of the C indexing design feature encapsulated earlier in the generation process. Recall that that design feature maps the domain language indexing system for neighborhoods (i.e., [-n, +n]) to a C language style of indexing (i.e., [0, 2n]).

VI. THE DSLGEN™ PROTOTYPE

DSLGen™ is the culmination of a six year R&D effort and comprises about 52KLOC of CommonLisp and CLOS running on Franz Allegro, version 8.2. A key component of the architecture upon which many of the other components are built is a general pattern matching system with backtracking, which is built using continuations. Built on top of the pattern matcher is a transformation system that includes several flavors of transformations (e.g., general, MTs, generic components, and deferred, which are used to move newly created subtrees up the abstract syntax tree). The transformations are generator phase specific (i.e., they are only enabled during named generator phases, e.g., the SyntheticDesign phase).

The partial evaluator and several specialized inference subsystems are also heavy users of the pattern matcher. The inference systems include a type inference system, a backchaining rule system used for inferring PL based loops

from the logical loop constraints, a logical expression simplifier based on logical subsumption and an inequality inference engine based on Fourier-Motzkin elimination, which is used for inferring relationships among logical architecture elements (e.g., inferring that $(i == (m - 1))$ is false when $(i == 0)$ and $(m > 1)$ are true).

As to generation times, the Sobel example on RGB images with partitioning but without load leveling or threads takes about 75 seconds to generate an Abstract Syntax Tree (AST) for C (or 40 to 50 seconds if not generating history and traces). Adding in the surface syntax to generate the text-based C files adds an additional 15 to 20 seconds. Adding multicore with threads, SIMD and various other architectural complexities increases generation times by small, linear amounts.

VII. PERFORMANCE TESTING

Generated code was tested with a selection of implementation variations on a 4 core, 3.33 GHz Velocity brand computer with 12 GB of real and 24 GB of virtual memory. The computer is built on the Intel i7 CPU with Turbo mode, which allows overclocking when the CPU is running under maximum temperature and power specification. It has 8 virtual processors. The code was compiled with Microsoft's Visual Studio 2008 C/C++ compiler.

The test data was a 215 by 215 pixel image in RGB format with a 24 bit pixel depth. The chosen computations included Sobel and Wallis edge detection methods [27] since they put a greater computational load on the machine than other possible computations might. In addition, Sobel provides one of the more serious challenges to the generator in that it requires use of virtually all of the generation facilities. The testing also included image Average and Unsharp Mask [27] (often used to sharpen Mammogram images), both of which have lighter computational loads.

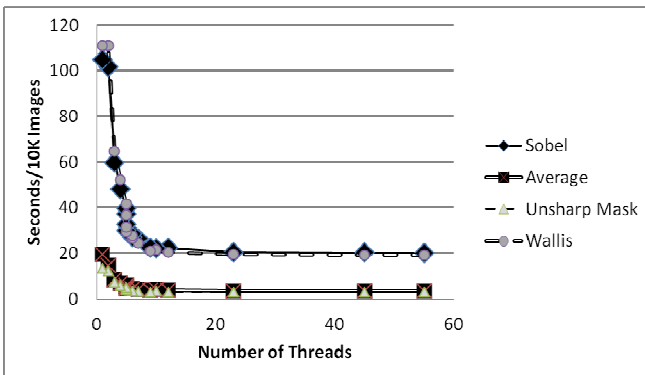


Figure 5. Performance vs. thread count

Figure 5 shows the results for various computations decomposed into threads to be run in parallel. These tests were run 10,000 times per image. For Sobel, the best performance was achieved at 55 threads, which required approximately 20.3 seconds to run the full set, or about 2

milliseconds per image. The worst results were with two threads, one for the edge cases and one for the center, which required approximately 105 seconds for the 10,000 images or a bit over a 10 millisecond per image. This was roughly the same time required for the calibration case, a hand coded version compiled to use no parallelism of any kind. Notice that the time drops quickly with five threads (i.e., one for the edges and four for the image center), taking about 32.8 seconds for the full set of images or about 3.3 milliseconds per image. This is about what simple logic would expect with four cores. However, the time continues to improve modestly for each five or so additional threads until it begins to level out at about 20.5 seconds at about 23 threads. Thereafter, the improvement is a tenth of a second or so for five or so additional threads. It is somewhat counter intuitive that one should get any improvement at all after the image has been evenly decomposed over the four cores. It is not entirely clear why this occurs but our current hypothesis is that it may be the "GPU effect" where many threads can mask memory, cache or other kind of latency if thread switching is efficient enough. Also, fast thread switching among virtual processors in the hardware (called Hyperthreading) may play a role. The target computer has two virtual processors per core and this is known to increase overall performance in many cases.

Other test cases with different kinds of image processing functions show similar behavior although the computational loads vary based on the nature of the computation. Sobel and Wallis have computational heavy loads that just simply require hefty computational capacity. Sobel employs square roots and Wallis uses logarithms. On the other hand, Average and Unsharp Mask are both light weight computations that employ little more than addition and division. Hence, they require less computational capacity as is clear from the graph.

With the addition of SIMD instructions, the added improvement ranges from about a 14% improvement for few or no threads to 36% for the maximum number of threads tested. With only two threads, Sobel took 87.2 seconds for all 10K images or 8.7 milliseconds per image, whereas with 55 threads, it took 12.87 seconds for all 10K images or about 1.3 milliseconds per image.

VIII. RELATED RESEARCH

A key difference between most previous research and DSLGenTM is that DSLGenTM starts working strictly in the problem domain and programming process domain rather than the PL domain. Virtually all previous research chooses representation systems that are based to some degree upon PL constructs or abstractions thereof. This includes compiling technology, generator technology [4] [5] [7] [21] [28] [29], computer aided software engineering (CASE) [13], model driven engineering [21], Aspect Oriented Programming (AOP) [17], Anticipatory Optimization Generation (AOG) [6] [7], general optimization based methods [1] [19] [16], parallel or specialty programming languages [8] [12], programming languages superficially similar to DSLGenTM's partitioning model [14], programming language augmentation systems [15] [26],

maintenance support systems [1] [2], refactoring [18] and other related technology and methods for creating implementation code from a specification of a computation. This representational choice forces conventional generation technologies to introduce design and PL forms, implementation structures, organizational commitments and other execution platform based details too early and thereby make design decisions about the architecture of the solution that will prevent other desired design decisions from being made later. Or at least, it will make those other desired design decisions require revision of the model or design and often difficult to automate.

In general, there are two important properties that differentiate these various approaches from DSLGenTM: 1) The specifications of the computations in these approaches are not invariant over a variety of execution platform architectures, and 2) target program implementations exploiting specific high capability features cannot be fully and automatically generated without compromising the invariance property. That is, user action is required either to revise the computational specification model to fit the new execution platform or to extend an overly abstract and therefore incomplete input specification to target a specific execution platform. Generation of target program implementations for a variety of execution platforms that exploit the execution platform features (e.g., multicore parallelism, vector instructions, etc.) requires human redesign or reprogramming in one form or another. For example, in these approaches, the transition from one execution architecture (e.g., simple Von Neumann) to another (e.g., multicore and/or vector machines) requires user action to adapt the computation specification or model to the new execution architecture.

In many cases, these conventional technologies often force a top down, reductionist approach to design where the top level programming structure and the essence of its algorithm are expressed first and then the constituent essence is recursively extended step by step until the lowest level of PL details are expressed. However, that initial structure may be incompatible with some desired design requirements or features that are addressed later in the development or generation process. The initial design may have to be reorganized to introduce such design requirements or features. For example, the requirement to fully exploit a multicore computer requires a significant, difficult and many step reorganization to fully exploit the performance improvements possible with multicore. Automation of such reorganizations at the programming language level is seriously complicated and except for relatively simple cases is prone to failure. This is why compilers that can compile programs written and optimized for one execution platform are often unable to satisfactorily compile the same programs for a different execution platform with an architecture that employs a significantly different model for high capability execution and fully exploit the high capability features of the new architecture. For example, programs written for the pre-2000 era Intel platforms are largely unable to be automatically translated to fully exploit the multicore parallelism of the more recent Intel platforms. Human based

reprogramming is almost always necessary to fully exploit the multicore parallelism.

While much research has been highly PL oriented, some research is clearly working in the problem domain. A prime example is the work of Jim Neighbors, who introduced the idea of using domain specific information in program generation. [23] [24] [25] His approach is to map from purely problem domain oriented languages through a series of language to language mappings, incrementally evolving to pure programming language representations. While DSLGenTM is consistent with that spirit, the underlying machinery (e.g., the non-top-down design approach, the non-PL logical architecture model, the APCs, the incremental design feature encapsulation and the incremental addition of sets of PL features phase by phase) distinguishes the DSLGenTM approach from Neighbor's work. Nevertheless, Neighbors' work has made significant contributions to program generation from which this work has benefited.

IX. CONTRIBUTIONS

The contributions of this work are due in large part to the fact that this work breaks with convention in a number of ways. Perhaps the most important break is avoiding the PL domain in the initial modeling process. This allows the *implementation neutrality* of the *INS* and allows *the separation of the INS from the specification of the execution platform (EXPS)* while still allowing the generated programs to *exploit the full range of high capability features* of the EXPS. While some systems emphasize language neutrality [29] rather than implementation neutrality, their specifications clearly derive from the PL domain and they therefore inherit the liabilities of the PL domain.

The ability of DSLGenTM to exploit high capability features arises from another important contribution, specifically, the design representation system based on *associative programming constraints*. The design representation system allows the initial and early stage designs to be organized as *logical architectures* thereby allowing the system to *operate in the problem and programming process domains* and to introduce PL constructions and assumptions incrementally. Operating with *problem domain concepts* such as edge and center partitions allows DSLGenTM to begin to manipulate and extend the LA without (initially) being restricted by the constraints inherent to programming languages.

Organizing the *IL definitions* as *provisional transformations that are malleable* provides the opportunity for *incrementally adding design features* by using higher order transformations to revise the IL definitions to incorporate those features but still defer casting them into programming language constructs until late in the generation process. Thus, the IL becomes the stand-in or precursor representation for the code details that have yet to be concretely determined. For example, expressions like Partestx(sp) can stand-in for code or meta-information (e.g., assertions) that cannot be refined to concrete form until the implementation context (i.e., a specific partition) and locale (i.e., the location in the AST) are concretely and finally

determined. And when that context is eventually pinned down (e.g., to Edge1), Partext(sp) can be specialized (e.g., to Partext(sp-Edge1)), which will move it a step closer to refinement into a concrete logical expression.

DSLGenTM relies heavily on *inference and implication*. For example, the APCs are described by a set of logical assertions that are augmented as the design progresses. This allows *architectural features and programming clichés to be expressed inferentially* rather than structurally and proscriptively. This defers making PL level design decisions. These PL representational forms are hard to revise, change and manipulate. For example, in DSLGenTM, adding messy design details and programming clichés can be deferred until the broad architectural structure is settled.

In summary, DSLGenTM represents a fundamentally new paradigm for program generation.

X. ACKNOWLEDGMENTS

I want to thank two contributors who recently began helping with the commercialization of DSLGenTM – Mitch Lubars and Rob Pettengill. Mitch did the performance testing and implemented the Fourier-Motzkin based inference engine. Additionally, Mitch wrote a prototype of a Slicer/Slicee target program from which the author abstracted the Slicer/Slicee design pattern used to support one class of synthetic partitions. Rob is working on a search and bookmarking facility for the transformation history debugger, which is used by Domain Engineers who intend to extend and customize DSLGenTM's domains and transformations. Additionally, both Mitch and Rob read this paper and provided many comments and suggestions that improved it.

References

- [1] Robert L. Akers, Ira D. Baxter, Michael Mehlich, Brian J. Ellis, and Kenn R. Luecke, "Case study: Re-engineering C++ component models via automatic program transformation," *Information and Software Technology* 49, 2007, pp. 275-291.
- [2] Ira D. Baxter, Christopher Pidgeon, and Michael Mehlich, "DMS[®]: Program Transformations for practical scalable software evolution," *International Conference of Software Engineering*, May 2004, pp. 10.
- [3] David F. Bacon, Susan L. Graham, and Oliver J. Sharp, "Compiler transformations for high-performance computing," *ACM Surveys*, Vol. 26, No. 4, December, 1994, pp. 345-420.
- [4] Don Batory, Vivek Singhal, Marty Sirkin, and Jeff Thomas, "Scalable software libraries," *Symposium on the Foundations of Software Engineering*. Los Angeles, California, 1993, pp. 191-199.
- [5] Ted J. Biggerstaff, "A perspective of generative reuse, annals of software engineering," *Baltzer Science Publishers*, AE Bussum, The Netherlands, 1998, pp.169-226.
- [6] Ted J. Biggerstaff, "Fixing some transformation problems" *Automated Software Engineering Conference*, Cocoa Beach, Florida, 1999, pp. 10.
- [7] Ted J. Biggerstaff, "A new architecture of transformation-based generators," *IEEE Transactions on Software Engineering*, Vol. 30, No. 12, Dec., 2004, 1036-1054.
- [8] Ted J. Biggerstaff, "Automated partitioning of a computation for parallel or other high capability architecture," Patent no. 8,060,857, United States Patent and Trademark Office, filed January 31, 2009, issued November 15, 2011.
- [9] Ted J. Biggerstaff, "Non-localized constraints for automated program generation," United States Patent and Trademark Office, Patent no. 8,225,277, filed April 25, 2010, issued July 17, 2012.
- [10] Ted J. Biggerstaff, "Synthetic partitioning for imposing implementation design patterns onto logical architectures of computations," United States Patent and Trademark Office, Patent no. 8,327,321, filed August 27, 2011, issued Dec. 4, 2012.
- [11] Guy E. Blelloch, Jonathan C. Hardwick, Siddhartha Chatterjee, Jay Sipelstein, and Marco Zagha, "Implementation of a portable nested data-parallel language," in *Proceedings of PPOPP '93 Proceedings of the fourth ACM SIGPLAN symposium on Principles and practice of parallel programming*, 1993, 102-111
- [12] Guy Blelloch, "Programming parallel algorithms," *Communications of the ACM*, 39 (3), March, 1996, pp. 85-97.
- [13] CASE Tools, See http://en.wikipedia.org/wiki/Rational_Rose .
- [14] Bradford L. Chamberlain, Choi, Sung-Eun, Deitz, Steven J. and Snyder, Lawrence, "The high-level parallel language ZPL improves productivity and performance," *Proceedings of the IEEE International Workshop on Productivity and Performance in High-End Computing*, 2004, pp. 1-10.
- [15] Barbara Chapman, Gabriele Jost and Ruud Van Der Pas, *Using OpenMP*, MIT Press, 2008.
- [16] Daniel E. Cooke, J. Nelson Rushton, Brad Nemanich, Robert G. Watson, and Per Andersen, "Normalize, transpose, and distribute: an automatic approach to handling nonscalars," *ACM Transactions on Programming Languages and Systems*, Vol. 30, No. 2, 2008, pp. 49.
- [17] Tzilla Elrad, Robert E. Filman and Atef Bader, "Aspect-oriented programming," *Communications of the ACM*, Vol. 44, No. 10, 2001, pp. 29-32.
- [18] Martin Fowler, Kent Beck, John Brant and William Opdyke, "Improving the design of existing code by refactoring," Addison-Wesley, 2000, pp 431.
- [19] M. W. Hall, S. P. Amarasinghe, B. R. Murphy, S. W. Liao, and M. S. Lam, "Interprocedural parallelization analysis in SUIF," *ACM Transactions on Programming Languages and Systems*, Vol. 27, No. 4, July, 2005, pp. 662-731.
- [20] Neil D. Jones, "An introduction to partial evaluation," *ACM Computing Surveys*, Vol. 28, No. 3, 1996, pp. 480-503.
- [21] Steve Macdonald, Kai Tan, Jonathan Schaeffer, and Duane Szafron, "Deferring design pattern decisions and automating structural pattern changes using a design-pattern-based programming system," *ACM Transactions on Programming Languages and Systems*, Vol 31, No. 3, April, 2009.
- [22] Model Driven Engineering. See http://en.wikipedia.org/wiki/Model-driven_engineering .
- [23] James M. Neighbors, "The Draco approach to constructing software from reusable components," *IEEE Transactions on Software Engineering*, SE-10 (5), (Sept. 1984) pp 564-573.
- [24] James M. Neighbors, "Draco: a method for engineering reusable software systems," In: Ted Biggerstaff and Alan Perlis (eds.): *Software Reusability*, Addison-Wesley/ACM Press (1989), pp. 295-319
- [25] James M. Neighbors, see <http://www.bayfronttechnologies.com/>.
- [26] OpenMP Architecture Review Board, "OpenMP Application Program Interface," Version 3.0, May 2008.
- [27] Gerhard X. Ritter and Joseph N. Wilson, *The Handbook of Computer Vision Algorithms in Image Algebra*, CRC Press, 1996.
- [28] Kai Tan, Duane Szafron, Jonathan Schaeffer, John Anvik And Steve Macdonald, "Using generative design patterns to generate parallel code for a distributed memory environment," *Proceedings of the ninth ACM SIGPLAN symposium on Principles and practice of parallel programming*, June, 2003, pp. 203-215.
- [29] Satnam Singh, "Computing without processors," *CACM*, Sept. 2011, pp. 46-54.

