Creating Explicit Communication in SoC Models Using Interactive Re-Coding*

Pramod Chandraiah,   Junyu Peng,   Rainer Dömer

Center for Embedded Computer Systems
University of California, Irvine
California, USA
{pramodc, pengj, doemer}@cecs.uci.edu

Abstract— Communication exploration has become a critical step during SoC design. Researchers in the CAD community have proposed fast and efficient techniques for comprehensive design space exploration to expedite this critical design step. Although these advances have been helpful in reducing the design time significantly, the overall design time of the system is still a bottleneck. All these techniques assume the availability of an initial SoC input model with explicit communication, whose quality significantly impacts the effectiveness of the communication exploration techniques. Today, these initial models need to be manually written by engineers, which is tedious, error-prone and time consuming. In fact, our studies on industrial-size examples have shown that about 50% of the communication exploration time is spent on coding and re-coding of the initial specification model. In this paper, we propose an efficient interactive approach to explicit communication creation by automating some of the common coding tasks in specification models for communication exploration. Our results show significant savings in designer time.

Fig. 1. Motivation: Extent of automation in refinement-based design flow.

I. INTRODUCTION

The increasing popularity of MPSoC architectures has made the task of communication exploration more significant than ever. Researchers have proposed different communication exploration techniques to facilitate early determination of communication architecture and configurations. By conducting the exploration at higher abstraction levels, such as TLM [4], CCATB [10], it is possible to explore, evaluate and verify different communication architectures early in the design cycle. All these technological advances along with the progresses in the area of code profiling, code refinement and accurate performance and power estimation have significantly reduced the overall development time of embedded systems. However, design time is still a bottleneck in the production of systems, and further reduction through automation is necessary. One critical aspect neglected in optimization efforts so far is the design specification phase, where the intended design is captured and modeled for use in the design flow.

Each design flow expects a specific type of input model. This model needs to be either hand-written from scratch, or adapted from an initial reference model. While much of the research has 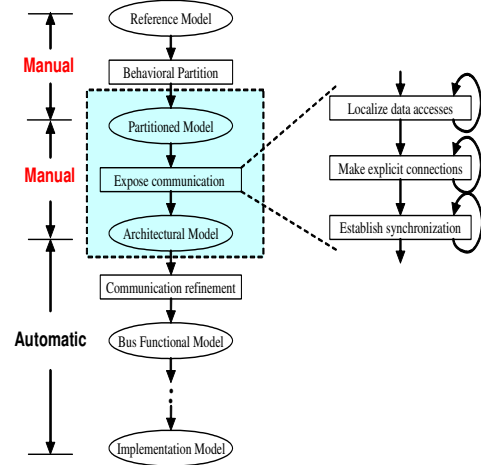focused on SoC synthesis and refinement tools, little has been done to support the designer in forming these input models.

A. Motivation

Most of the existing communication exploration design flows start with an initial executable model with a communication structure as supported by the tool. These models are written by designers who need to understand the requirements and the limitations of the tool. Given a suitable initial model, the communication exploration is usually conducted automatically using successive refinement steps. However, often there is little or no tool support in creating the initial input model. Since the quality of the communication exploration depends largely on the quality of this model, designers have to invest considerable time and effort in coding and re-coding the model.

A.1   Study of MP3 Decoder Design

In order to study the intricacies and complications involved in writing a proper system specification model, we have applied a top-down design methodology, shown in Figure 1, to the example of a MP3 audio decoder, an industry-size application. The figure shows the complete design flow starting from an abstract reference model down to the implementation model. In this paper, we only focus on the refinement steps that were necessary

to generate the Architectural Model for communication exploration.

The communication exploration is undertaken after different hardware-software partitions in the design are identified and mapped onto respective processing elements. During communication exploration, different bus architectures and configurations are then evaluated and the most suitable communication alternative, meeting power and performance requirement, is finalized. The other key task involved during this step is the communication refinement, during which the decisions made by the designer are implemented to generate an executable model. As shown in Figure 1, the task of communication exploration was automated, to the extent that model generation is fully automatic, and the designer only makes the design decisions such as bus allocation, mapping and scheduling. Due to this automation, starting from a Partitioned Model [1], we were able to explore different communication alternatives [1] and implement the bus-functional model of the MP3 decoder in less than 1 day. However, specifying proper communication structure in the initial Partitioned Model was the main bottleneck of the whole process. For our MP3 design example, more than 75% of the overall communication exploration time was spent on creating this model. Also, we need to emphasize that capturing this model is not a one time task. Every time a change in the design is required for a successful refinement step, it is necessary to *re-code*/change the input specification, making the whole task of coding model iterative. Such interruptions in the design flow cause costly delays. The importance of this *re-coding* effort is also emphasized in [6, 3].

In conclusion, any step toward automation of model coding and re-coding is highly desirable and will likely improve overall design time significantly. In Section II, we discuss specification modeling goals, and a modeling example, followed by related work. In Section III, we present our solution to the modeling problem, an interactive source re-coder. In Section IV we will present the refinement tasks in our source re-coder that simplify and speed-up the creation of a proper Architecture Model. Section V lists our experimental results. Finally, we will draw conclusions and outline future work in Section VI.

II. MODEL RE-CODING

To address the bottleneck of model re-coding, we will now first describe the significance of having a good model for communication exploration, and then analyze the various steps of this re-coding task by use of a simple example.

A. Significance of Explicit Communication in SoC Specification

To enable fast and efficient communication exploration, a proper communication structure in the initial specification is necessary for the refinement tools to function and to be effective. Separation of the computation from the communication aspects in the specification makes automatic communication refinement and exploration possible [2]. Figure 2 shows the

---

---

[1] Note that the generation of Partitioned Model from Reference Model is a separate problem and is not discussed in this paper.
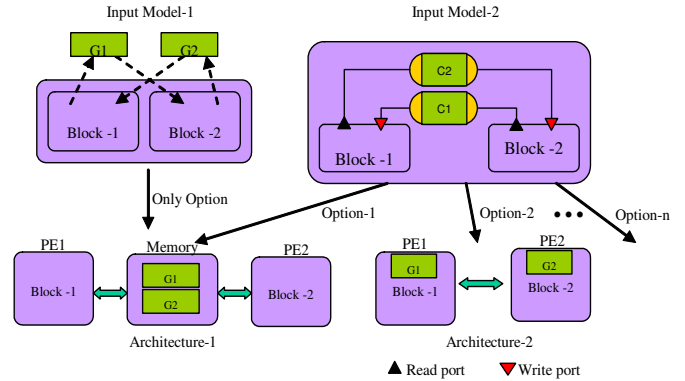

Fig. 2. Significance of explicit SoC communication.

benefit of having a proper model by means of a simplified example. Input Model-1 has the computation interleaved with the communication, making analysis difficult for automatic exploration tools. Architecture-1 is the only possible implementation through automation. In contrast, the Input Model-2 captures all the communication between the design partitions explicitly using synchronizing channels, thus enabling multiple explorations alternatives, Architecture-1, Architecture-2 and more. Global variables in Partitioned Models, such as the first model in Figure 2, are commonly used, especially when the Partitioned model itself is re-coded from a C like reference model. Global variables hide the communication between functions in a program because they do not appear as function parameters or return results of functions. Since they become globally available to all the functions in the program, programmers use this feature for convenience. However, a good specification model requires the communication modeled explicitly, separated from the computation. So, the hidden communication through global variables must be explicitly exposed.

B. Re-coding to Expose Communication

Figure 3 shows the typical refinement steps necessary to re-code a Partitioned Model into a properly structured SoC Architecture Model. We start with a Partitioned Model, in which each partition represents the tasks running on different Processing Elements (PEs). Each partition might in turn contain more behaviors within them. The partitions communicate with each other using global variables just like the way its ancestor C reference model did. In the first step, we reduce the number of such global variables by localizing them. During this optimization step, any global variable whose usage is restricted to just one of the behavior blocks is moved into that leaf behavior in the partition. This operation does not introduce any other changes and applies only to variables whose access is restricted to just one behavior.

In the second step, the implicit access of data through the use of the remaining global variables is replaced by making actual connections into each partitions. That is, the global variables that are shared across the multiple partitions or behaviors are now made accessible only through ports. This means that every access to the global variable is replaced with an access to a corresponding port in all levels of structural hierarchy in
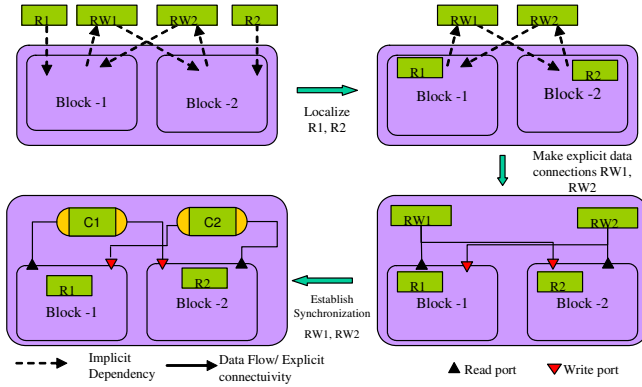
Fig. 3. Re-Coding Partitioned Model into an Architecture Model.

Fig. 4. Conceptual Structure of the Source Re-Coder.

each partition. If necessary, the changes are also propagated into functions by introducing new function parameters and arguments. The original global variable is deleted.

In the third step, the access to the variables shared across each of these partitions are synchronized by replacing the variables with communication channels. All the accesses to variables in each partition are replaced with use of an appropriate interface function (*Send/Receive* or *read/write*) implemented by the channel. The resulting Architecture Model after these transformations is then suitable for communication exploration because now the communication across partitions is explicitly captured using synchronizing channels and ports which represent future bus. Since the communication to be mapped onto the buses is completely exposed, design tools are free to explore unlimited alternatives. Most of the steps performed during communication exploration can now be automated.

C. Automated Re-Coding

Apart from the time consuming mundane textual operations, re-coding a model involves a lot of decision making. Many of these decisions can only be taken by the designer. Since the model being generated depends on the tool it will be input to, only the designer can decide on its structure. Apart from this reason, there are some scenarios where complete automation cannot be possible. For example, while determining the access information of global variables, the designer has to deal with pointers. Since pointer analysis [5] is a problem lingering for many years, user interaction becomes necessary to generate efficient and optimized models. Thus, decisions need to be taken by the designer, but the tedious recoding of the model can be automated.

To leverage the idea of effective automation, we need to distinguish re-coding tasks that can be automated from decision tasks that require designer's knowledge. Textual re-coding operations such as changing scope of variables, replacing variables and their access with channel equivalents, changing variable types, introducing/deleting ports through the hierarchy, can be automated if the decision is made by the designer. By such an automation, the designer is relieved from mundane text editing tasks and can focus on actual modeling decisions.

D. Related Work

Many of the system synthesis design flows proposed by researchers in the past provide little or no support in modeling the initial system Architecture Model. The designer has to invest significant time to create this specification. For example, [12] presents a methodology to automatically generate models and implementations of network-oriented SoC models. But the design flow starts from an architecture model which needs to be provided by the designer. [10] proposes a higher modeling abstraction for bus-based communication exploration, but the creation of the initial model of acceptable quality for exploration has to be performed by the designer. [8] focuses on automating mapping of communication between system components onto a communication architecture template, but requires an initial system specification partitioned into hardware and software from the designer.

Through our transformations we provide faster and efficient means of creating a suitable Architecture model. Our transformations are interactive and give the designer complete control ("designer-in-the-loop") to code/re-code the model in order to arrive at the most suitable design implementation.

Further, the idea of interactive transformations has been employed successfully in the parallel programming paradigm. The ParaScope editor [7] for Fortran and SUIF explorer [9] for C/Fortran provide program transformations to parallelize a program relieve the programmer from tedious manual typing. SUIF explorer, which is based on the SUIF compiler infrastructure, provides graphical means of setting compiler directives, but does not support editing. ParaScope provides the user with powerful interactive program transformations and reconstructs the dependency information, incrementally, while editing. Of all, ParaScope combines the most features. Our goal is to build similar and more advanced capabilities into our source re-coder, aiming at advanced features for analysis and transformations for efficient creation of effective system specifications in SLDL. Our work will augment the existing design flows by providing further automation.

III. INTERACTIVE SOURCE RE-CODER

To aid the designer in coding and re-coding, we propose a source *re-coder*. Our source re-coder is a controlled, interactive approach to implement analysis and refinement tasks. In other words, it is an intelligent union of editor, compiler, and powerful transformation and analysis tools. Unlike other program transformation tools, our re-coder keeps the designer in the loop and provides complete control to generate and modify a model suitable for her/his design flow. By making the re-

coding process interactive, we rely on the designer to concur, augment or overrule the analysis results of the tool, and use the combined intelligence of the re-coder and the designer for the modeling tasks.

Our re-coder supports re-modeling of SLDL models at all levels of abstraction. It can be used to re-code intermediate design models as well as the reference C implementation to generate the initial specification model. The conceptual structure of our source re-coder is shown in Figure 4. It consists of 5 main components:

- A textual editor maintaining the textual document object
- An Abstract Syntax Tree (AST) of the design model
- Preprocessor and Parser to convert the document object into AST
- Transformation and analysis tool set
- Code generator to apply changes in the AST to the document object

A QT [13] and Scintilla [11] based textual editor is the frontend of our source re-coder. The editor is adapted to provide basic features like syntax highlighting, auto-completion, search, ctags, text folding, bookmarks, undo-redo, and more, for programming languages including C and C++, and SLDLs SpecC [2] and SystemC.

The AST [14], maintains a comprehensive data structure of the design needed for analyzing and transforming the program and also provides a set of operations on each object. This AST is created by a parser when the file is initially loaded and subsequently is incrementally maintained with designer's actions.

The transformation and analysis tool set is the heart of our source re-coder. All re-coding tasks invoked by the user are implemented by these refinement tools. When the designer points to an object in the source window, a node corresponding to the pointed co-ordinates is located in the AST, and a list of available and possible operations are provided in a context menu. For example, when the designer points to a global variable, then the list of possible transformations includes changing its scope, renaming, deleting, and finding dependents. A SLDL source code generator provides the necessary update of text in the editor when the modifications are made to the AST by the transformation tools.

IV. EXPOSING COMMUNICATION

We will now discuss specific transformations necessary to create an Architecture Model suitable for communication exploration, as outlined in Figure 3. Exposing communication means to make the communication between different hardware software partitions explicit. This involves 3 main tasks, Localizing global accesses, establishing explicit connectivity and introducing synchronization. These steps are discussed in detail in the following sections.

A. Localize Global Accesses

In the first step, global variables, whose usage is restricted to just one of the partitions, are made members of that partition. Such scenarios often occur in partitioned models obtained from a C reference implementation. Localizing such variables will eliminate the unnecessary communication between the partitions and also makes the analysis of the specification easier which enables further optimization. Figure 5(a) shows a simple initial model in SpecC SLDL. The modified model after localizing 2 of the global variables is shown in Figure 5(b). The variables *R1* and *R2* are made members of the blocks in which they are used. The procedure implementing this transformation is given below.

- From the current cursor position, locate the global variable pointed to by the designer
- Find the list of behaviors and functions accessing this global variable.
- Move this variable into the only behavior accessing it and delete the global variable

B. Explicit Connectivity

To enable automatic communication refinement, all the communication between partitions need to be explicitly specified. Global variables pose problems and need to be localized and accessed through ports. During this step, the global variables are removed from the global scope and moved into the lowest parent behavior containing the accessors. Following this, each partition accessing the global variable is modified to access the global variable through a port. If necessary, the access is routed through the entire structural hierarchy introducing new ports, port maps, function arguments and more in each partition.

Figure 5(c) shows the modified model after changing the scope of the global variables *RW1* and *RW2* and also the resulting new ports. The procedure implementing this transformation is outlined below.

- Obtain the global variable at the current cursor position
- Find the lowest common parent behavior containing the accesses and move the variable to that level
- Provide access to the variable by recursively inserting ports in all the behaviors accessing this variable
- Delete the original global variable

This operation is not restricted to just the global variables. Depending on the model being derived, designer might want to make some variables global or move a variable to higher hierarchical levels. This requires expanding the scope of a variable. The procedure for such an operation is based on similar lines.

C. Introduce Synchronization

The accesses to variables that are shared between concurrent partitions need to be synchronized. Such accesses are replaced with channels of appropriate protocols to provide necessary synchronization. This change in turn requires all the behaviors and the functions accessing the replaced variable to be modified to access the variable through the channel interface. Figure 5(d) shows the resulting model with the inserted channel calls. The reading and writing of variables is replaced with the appropriate interface function of the channel. The procedure implementing this transformation is given below.

- Obtain the variable at the current cursor position

```
/* Global variables */
int R1, R2;
int RW1, RW2;

/*Top level behavior */
behavior Main( ) {
  int var1, var2, var3;

  b1 B1(var1, var2);
  b2 B2(var2, var3);

  int main(void) {
    B1.main();
    B2.main();
  }
};
/* Sub modules */
behavior b1(in int i1, out int o1) {
  void main (void) {
    o1 = R1*RW2*i1;
    if(RW2) RW1 = ((R1*RW2)*i1)&1;

  }
};
behavior b2(in int i1, out int o1) {
  void main(void)  {
    o1 = R2*RW1*i1;
    if(RW1) RW2 = ((R2*RW1)*i1)&1;
  }
};
```

(a) Model-1: Initial Model

```
/* Global variables */
int RW1, RW2;

/*Top level behavior */
behavior Main( ) {
  int var1, var2, var3;

  b1 B1(var1, var2);
  b2 B2(var2, var3);

  int main(void) {
    B1.main();
    B2.main();
  }
};

/* Sub modules */
behavior b1(in int i1, out int o1) {
  int R1;
  void main (void)  {
    o1 = R1*RW2*i1;
    if(RW2) RW1 = ((R1*RW2)*i1)&1;
  }
};
behavior b2(in int i1, out int o1) {
  int R2;
  void main (void)  {
    o1 = R2*RW1*i1;
    if(RW1) RW2 = ((R2*RW1)*i1)&1;
  }
};
```

(b) Model-2: After Localization

```
/*Top level behavior */
behavior Main() {
  int var1, var2, var3;
  int RW1, RW2; /* Now moved here, no longer global*/
  b1 B1(var1, var2, RW1, RW2);
  b2 B2(var2, var3, RW2, RW1);
  int main(void) {
    B1.main();
    B2.main();
  }
};
/* No more Global variables */
behavior b1(in int i1, out int o1,
            out int RW1, in int RW2){
  int R1;
  void main (void)  {
    o1 = R1*RW2*i1;
    if(RW2) RW1 = ((R1*RW2)*i1)&1;
  }
};
behavior b2(in int i1, out int o1,
            out int RW2, in int RW1) {
  int R2;
  void main (void)  {
    o1 = R2*RW1*i1;
    if(RW1) RW2 = ((R2*RW1)*i1)&1;
  }
};
```

(c) Model-3: Explicit connectivity

```
/*Top level behavior */
behavior Main()
{
  int var1, var2, var3;

  c_fifo ch1; /*Channels instead of variables */
  c_fifo ch2;

  b1 B1(var1, var2, ch1, ch2);
  b2 B2(var2, var3, ch2, ch1);

  int main(void)
  {
    B1.main();
    B2.main();
  }
};
behavior b1(in int i1, out int o1, i_sender ch1,
            i_receiver ch2)  {
  int R1;
  int RW1; /*local variables*/
  void main (void)  {
    o1 = R1*(ch2.receive(sizeof(RW2)))*i1;
    if(RW2) RW1 = ((R1*RW2)*i1)&1;
    ch1.send(RW1);

  }
};
behavior b2(in int i1, out int o1, i_sender  ch2,
            i_receiver ch1) {
  int R2;
  int RW2;
  void main (void)  {
    o1 = R2*(ch2.receive(sizeof(RW1)))*i1;
    if(RW1) RW2 = ((R2*RW1)*i1)&1;
    ch2.send(sizeof(RW2));
  }
};
```
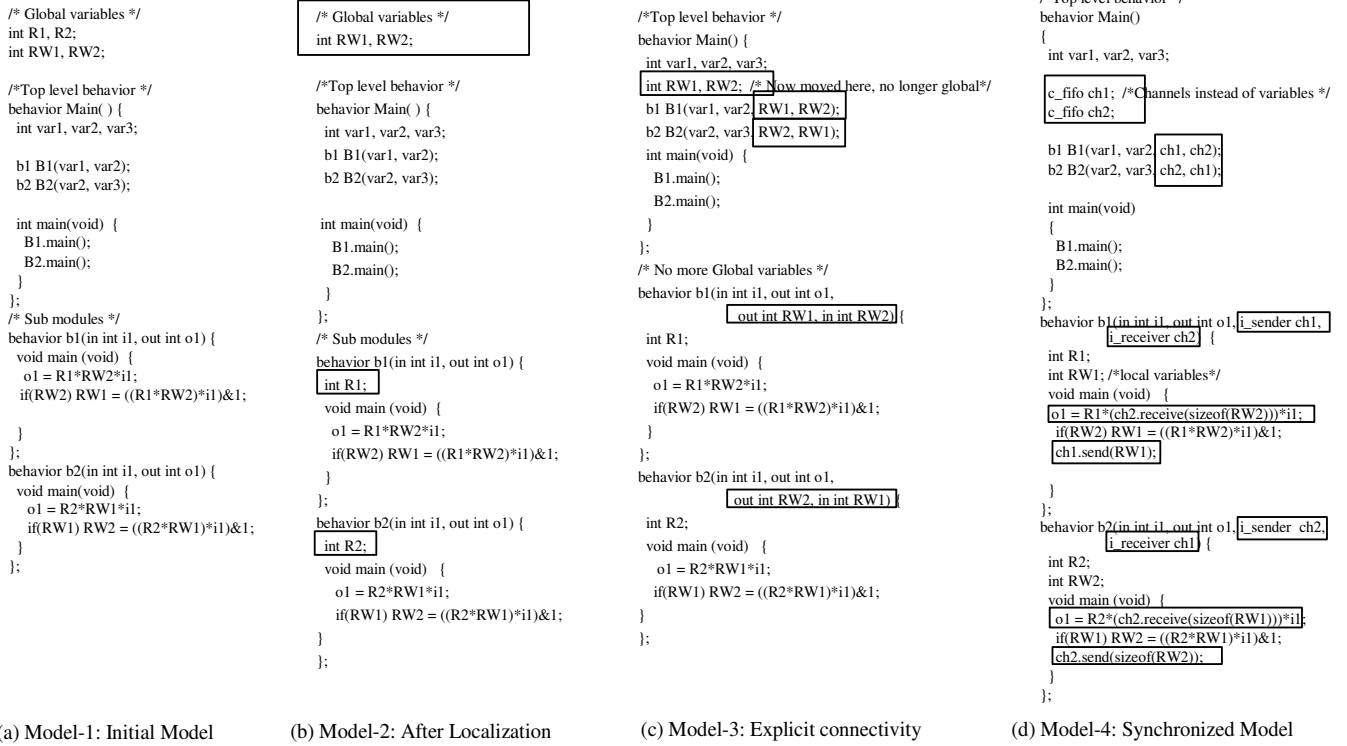
(d) Model-4: Synchronized Model

Fig. 5. Re-coding transformations on an example design model.

- Create a typed synchronization channel based on the type of the variable
- Change the type of ports of all the behaviors which were mapped to the original variable to the type of the channel's interface function
- Modify each access to the original variable to use the appropriate interface function of the channel

V. EXPERIMENTS AND RESULTS

We have implemented our source re-coder to support the SpecC SLDL and the C programming language. Additional support for SystemC is still in progress. Using our re-coder, the modeling can start either from C or SpecC. Currently, our re-coder provides complete support to derive an architecture model from a partitioned model. Apart from these transformations, the re-coder also implements scope-sensitive analysis tools to find variable dependencies to aid the designer in program comprehension. The transformation functions makes use of this analysis function for re-coding.

We will now assess the productivity gains resulting from our source re-coder. We have applied our source re-coder to a few real-life design examples to create architecture models and measured the productivity gains over deriving the same models manually. We considered 3 different examples, MP3 decoder, JPEG encoder, and GSM vocoder. We have timed the individual refinement tasks for the MP3 example and the results are shown in Table I. We started with a partitioned MP3 model with 4 partitions and contained all the global communication between these partitions using the transformation tools part of

our source re-coder. Note that each of these transformations require multiple lines of code change, and these lines often are distributed over the entire source code. Table I also gives the number of variables changed by each refinement task and the number of line changes in the resulting model. As shown in the table, establishing explicit connectivity involves inserting new ports through the structural hierarchy in the design, resulting in more changes than localizing operation. The Interactive Re-coding time is obtained by implementing all the transformations using our source re-coder. The manual time is obtained by actually realizing the transformations manually for a set of 10 variables and extrapolating the results for all the variables. Clearly, the productivity gain achieved using our re-coder over implementing these transformations manually is in the order of hundreds.

Table II shows the similar productivity gains achieved for other examples. Also provided is the statistics about the number of variables that were affected and the number of new ports and channels that were introduced during the process. Using our source re-coder, the complex transformations can be realized instantly with a click of button. Thus, exposing communication can be achieved in the order of seconds instead of hours.

VI. SUMMARY AND CONCLUSIONS

Automatic communication exploration tools require proper input models with all the communication across the partitions exposed explicitly. Absence of a proper communication structure significantly limits the possible exploration, if not making it impossible. Today, designers invest significant design time

TABLE I
PRODUCTIVITY GAIN ON THE INDIVIDUAL OPERATIONS ON THE MP3
DESIGN EXAMPLE

Operation	Variable. changes	Lines changed	Interactive Re-coding time (secs)	Manual time for 10 Vars. (mins)	Estimated Manual time (mins)	Productivity factor
Localizing	26	330	90	1.7	17	11
Explicit Connectivity	38	839	126	31	310	147
Synchronization	6	172	30	17	170	340
Total	70	1341	246	49.7	497	121

TABLE II
PRODUCTIVITY GAIN FOR DIFFERENT EXAMPLES

Properties	JPEG	MP3	GSM
Global Variables localized	8	70	83
New Ports added	2	146	163
New Channels added	1	6	2
Re-coding time (secs)	27	246	260
Estimated Manual time (mins)	53	497	585
Productivity factor	117	121	135

to create explicit communication in the input model. Since this modeling task heavily relies on the designers' knowledge and experience, little or no tool support is typically available for specification coding and re-coding. Usually, designers have to edit the design model manually, using simple text-based editors, requiring significant effort and time in tedious coding.

In this paper, we have proposed a novel approach to design specification and modeling that is based on decision making by the designer ("designer-in-the-loop") and automation of model analysis and transformation tasks. Eliminating mundane and error-prone editing tasks, our approach utilizes the precious time and effort of the system designer efficiently.

In particular, we have introduced an interactive source re-coder which integrates compilation, analysis and transformation tools into a text-based editor, to assist the designer in modeling and re-modeling of SoC designs. Our source re-coder is fully text-, syntax-, and semantics-aware, enabling powerful model analysis and transformation operations instantly.

From our experience with the system synthesis of different design examples, in this paper we presented the various tasks necessary to build the communication structure in the specification that can be used for automated tools. We have designed effective interactive re-coding tasks along with other analysis tools to aid the designer in program comprehension, modification, and refinement. Our experimental results clearly demonstrate that our interactive approach is not only feasible, but also effective. Analysis results or transformed code are presented to the user instantaneously, relieving the designer from tedious coding. Moreover, we have demonstrated tremendous productivity gains through the reduction of modeling time.

For the future, we will extend our approach to automate also the creation of the partitioned model from a C like reference model. In the long term, we plan to include further analysis and transformation tasks, including coupling with system profiling and estimation tools.

REFERENCES

[1] P. Chandraiah and R. Dömer. Specification and design of an mp3 audio decoder. Technical Report CECS-TR-05-04, Center for Embedded Computer Systems, University of California, Irvine, May 2005.

[2] A. Gerstlauer, R. Dömer, J. Peng, and D. D. Gajski. *System Design: A Practical Guide with SpecC*. Kluwer Academic Publishers, 2001.

[3] A. Gerstlauer, S. Zhao, D. D. Gajski, and A. M. Horak. SpecC system-level design methodology applied to the design of a GSM vocoder. In *Proceedings of the Workshop of Synthesis and System Integration of Mixed Information Technologies*, Kyoto, Japan, April 2000.

[4] F. Ghenassia. *Transaction-Level Modeling with SystemC : TLM Concepts and Applications for Embedded Systems*. Springer-Verlag, 2006.

[5] M. Hind. Pointer analysis: Haven't we solved this problem yet? In *PASTE '01: Proceedings of the 2001 ACM SIGPLAN-SIGSOFT workshop on Program analysis for software tools and engineering*, 2001.

[6] A. Jerraya, H. Tenhunen, and W. Wolf. Guest editors' introduction: Multiprocessor systems-on-chips. *Computer*, 38(7):36–40, 2005.

[7] K. Kennedy, K. S. McKinley, and C.-W. Tseng. Analysis and transformation in the ParaScope Editor. In *ACM International Conference on Supercomputing*, Cologne, Germany, 1991.

[8] K. Lahiri, A. Raghunathan, and S. Dey. Efficient exploration of the soc communication architecture design space. In *ICCAD*, Piscataway, NJ, USA, 2000.

[9] S.-W. Liao, A. Diwan, R. P. B. Jr., A. M. Ghuloum, and M. S. Lam. SUIF explorer: An interactive and interprocedural parallelizer. In *Principles Practice of Parallel Programming*, pages 37–48, 1999.

[10] S. Pasricha, N. Dutt, and M. Ben-Romdhane. Fast exploration of bus-based on-chip communication architectures. In *CODES+ISSS*, New York, NY, USA, 2004.

[11] Scintilla source code editing component. http://www.scintilla.org.

[12] D. Shin, A. Gerstlauer, R. Dömer, and D. D. Gajski. Automatic network generation for system-on-chip communication design. In '*CODES+ISSS*, New York, NY, USA, 2005.

[13] Trolltech Inc. Qt application development framework. http://www.trolltech.com/products/qt/.

[14] I. Viskic and R. Dömer. A flexible, syntax independent representation (SIR) for system level design models. In *Proceedings of EuroMicro Conference on Digital System Design*, Dubrovnik, Croatia, August 2006.