

# Debugging and Verification of Parallel Systems — the *picoChip* Way!

Daniel TOWNER, Gajinder PANESAR, Andrew DULLER,  
Alan GRAY and Will ROBBINS  
*picoChip Designs Ltd., Bath, UK*

**Abstract.** This paper describes the methods that have been developed for debugging and verifying systems using devices from the *picoArray*<sup>TM</sup> family. In order to increase the computational ability of these devices the hardware debugging support has been kept to a minimum and the methods and tools described take this into account. An example of how some of these methods have been used to produce an 802.16 system is given. The important features of the new PC102 device are outlined.

## 1 Introduction

The wireless communications field is experiencing a period of major expansion. This is happening all over the world and is not dominated by any one region in particular. In a field in which different standards are fixed for different regions of the world, where standards are in a state of flux or even where no standards exist, it is very costly to enter the market with a custom ASIC solution. What is required is a scalable programmable solution, which can cater for most, if not all these areas. To this end *picoChip* created the *picoArray*<sup>TM</sup> and a rich toolset.

The *picoArray*<sup>TM</sup> is a tiled processor architecture in which hundreds of processors are connected together using a deterministic interconnect [1, 2]. The level of parallelism is relatively fine grained with each processor having a small amount of local memory. Each processor runs a single process in its own memory space and they use “signals” to synchronise and communicate. Multiple *picoArray*<sup>TM</sup> devices may be connected together to form systems containing thousands of processors using on-chip peripherals which effectively extend the on-chip bus structure.

In order to provide a commercially viable, massively parallel, scalable solution, *picoChip* has had to re-think methods of debug and verification in the following areas:

**scale:** depending upon the target, systems solutions may require moderate or massive computational power. To address this, many *picoArray*<sup>TM</sup> devices may be connected together<sup>1</sup>, creating systems containing thousands of processors.

**reduced non-essential hardware:** in order to produce a commercially viable chip, silicon area is best used for computation. Specialised hardware for non-compute must be justified and the emphasis should be on system-wide rather than processor centric debug. Therefore conventional processor support, such as register or memory trace mechanisms, are not as useful as they are in uni-processor systems. In the *picoArray*<sup>TM</sup>, hardware support for debug has been kept to a minimum in order to allow more processors to be fitted onto a single device.

---

<sup>1</sup>*picoChip* has boards with up to 16 devices.

**environment:** the picoArray™ devices are designed for use in embedded environments, where to keep costs to a minimum, there are relatively few inputs and outputs. In such an environment, the bandwidth available for debugging traffic is limited. picoChip provides relatively light weight access to the picoArray™ via a JTAG interface and the microprocessor interface.

**communication and synchronisation:** many parallel systems use special purpose libraries (e.g., Posix threads [3], Message Passing Interface [4]) or language support mechanisms (e.g., Java threads [5]) to handle communication and synchronisation of parallel processes. This is very difficult to justify in an embedded system where memory and system cost are areas of concern. The picoArray™ uses a deterministic interconnect fabric called the picoBus™. This behaves like a blocking, double-entry FIFO between processes, and is used for communication and synchronisation. No run-time arbitration of the picoBus™ is necessary, enabling the hardware to be made much simpler, and removing a possible source of bugs.

Conventional debuggers are usually designed for a single processor, and normally require a reasonable amount of hardware support. Even debuggers from vendors who claim to support multi-processor systems will require hardware support. More importantly however, conventional debuggers will not scale to systems consisting of hundreds or thousands of processors.

picoChip's approach to debug and verification has been to exploit the programming paradigm provided by the picoArray™ [1, 2] and to create a rich set of tools to aid system-wide debug and verification. Two main categories can be identified: language features and software tools.

The picoArray™ tools support an input language which is a combination of VHDL [6], ANSI/ISO C and assembly language. Individual processes are written in C and assembly, while structural VHDL is used to describe how processes are connected together using signals. Signals are strongly typed, and have specified bandwidths. They may be synchronous or asynchronous, point-to-point or point-to-multi-point. Processes are statically created by describing their number and type in the source files — no runtime creation of processes is possible. Thus, after a system has been compiled, the complete set of processes, and their connections is known, and the system will behave deterministically.

The most common software tool for debugging is the symbolic debugger, which allows the programmer's original source code to be displayed, along with the contents of source variables, on a per-process basis. The use of symbolic debugging is commonplace, so will not be considered further in this paper. The additional mechanisms that have been developed are as follows:

**design browser:** this allows both static and dynamic analysis of a user's design. The overall structure of a system can be graphically displayed in a number of ways. This allows the user to verify that the overall system has been connected in the way that was intended and to visualise and navigate through a complex design which may have been coded by many people. In addition, dynamic analysis can be performed and it can be used to visualise problems such as those associated with data throughput and deadlock.

**simulation:** the cycle accurate simulator allows all of the processor's internal state to be viewed, including aspects that the real hardware does not allow. Typically the user would start their development here and then migrate to hardware.

**scripting:** the debugger can be programmed using Tcl/Tk. This allows the user to build on top of the basic system provided by the standard debugger.

**traces:** track the contents of signals, variables, registers, and memory over time.

**probes:** a special type of process which can be inserted into a user’s design during debugging, to enable complex real-time analysis.

**traffic analysis:** extract the states of the various communications and indicate the reasons for any stalling.

**file I/O:** stream data in to and out of a system, using files.

**activity display:** condensed graphical display of the state of many processes in a system.

The remainder of this paper is structured as follows. Section 2 contains an overview of the picoArray™ devices. We then describe how the debug and verification tools might be used over the lifetime of a design, from initial component to integrated system. We then describe each of the types of debug mechanism in more detail and finish with a conclusion.

## 2 The picoArray™ Concept

### 2.1 The picoArray™ Architecture

picoChip’s latest device - PC102 is based around the picoArray™ tiled processor architecture in which over 300 processors (3-way VLIW, Harvard architecture with local memory), and 14 co-processors (Function Accelerator Units or FAU) are interconnected by a 32-bit picoBus and programmable switches.

The term Array Element (AE) is used to describe either processors or co-processors (i.e., there are 322 AEs in the array). There are three processor variants which share the same basic structure: Standard AE (STAN), Control AE (CTRL) and Memory AE (MEM). Memory configuration and the numbers of communications ports vary between AE types.

### 2.2 Inter-processor Communications

Within the picoArray™ core, AEs are organised in a two dimensional grid, and communicate over a network of 32-bit buses (the picoBus™) and programmable bus switches. AEs are connected to the picoBus™ by ports. The ports act as nodes on the picoBus™ and provide a simple interface to the bus based on “put” and “get” instructions in the instruction set.

The inter-processor communication protocol is based on a time division multiplexing (TDM) scheme, where data transfers between processor ports occur during time slots, scheduled automatically by the tools, and controlled using the bus switches. The bus switch programming and the scheduling of data transfers is fixed at compile time, and requires no run-time arbitration. Figure 1 shows an example in which the switches have been set to form two different signals between processors. Signals may be point-to-point, or point-to-multi-point. In the latter case, the data transfer will not take place until all the processor ports involved in the transfer are ready. The total internal data bandwidth for the signals is 3.3 Tera bits per second (322 processors x 2 buses x 32-bits x 160MHz clock).

The default signal transfer mode is synchronous; data is not transferred until both the sender and receiver ports are ready for the transfer. If either is ready before the other then the transfer will be retried during the next available time slot. If, during a *put* instruction no buffer space is available then the processor will sleep (hence reducing power consumption) until space becomes available. In the same way, if during a *get* instruction there is no data available in the buffers then the processor will also sleep. Using this protocol ensures that no data can be lost.




Figure 1: picoArray™ Interconnect

There is also an asynchronous signal mode where transfer of data is not handshaken and in consequence data can be lost by being overwritten in the buffers without being read.

### 2.3 External Communications

The picoArray™ has three methods of external communications. They are:

- Microprocessor Interface (MPI),
- Inter-picoArray™ Interface (IPI),
- Asynchronous Data Interface (ADI).

These can all be connected to the picoBus™ and can be accessed using signals. The MPI can be used to configure picoArray™ devices and can be used by debugging tools for input and output of relatively low bandwidth (2.5 Giga bits per second) information.

The IPI is used to connect picoArray™ devices together and can be viewed as a way of extending the picoBus™ across devices.

The ADI is used for exchanging data with high bandwidth (5 Giga bits per second) external asynchronous data streams.

Each device has a single MPI and four interfaces which can be configured as either an IPI or an ADI.

## 3 A Method for Design and Debug

This section goes through a typical process that is used to create a picoArray™ based application.

### 3.1 System Decomposition

Typically this is done by hierarchically breaking down the problem into components consisting of processes connected by signals. Experience has shown that components generally contain a few tens of processes, however the number of processes required does not have to be specified at this stage. The boundaries of these components will also have signals defined and will eventually be connected to other parts of the system. The user will use knowledge of the real-time system being developed to specify signal properties, such as maximum bandwidth and signal type. The properties can be checked during integration using signal assertions, which are described in section 3.3.

### 3.2 Component Coding

Two approaches can be taken, the choice being dictated by the complexity of the component.

For small components in which the division into AEs can be determined easily these AEs can be coded using C or ASM and connected using appropriate signals.

For larger components it may be preferable to initially produce a functional representation using C. This can be simulated even when the code size exceeds the memory for any AE and allows functional testing of this component prior to its division into individual AEs.

Whichever approach is used the code can be tested by creating test harnesses using FileIO to mimic the external components. The symbolic debugger and its attendant tools can be used to find bugs within the AEs.

The migration of the code to hardware is eased by the fact that the same FileIO test harnesses produced for simulation can be used for verification. This highlights a huge advantage of the picoChip approach since testing on hardware can be performed at a very early stage which means that components can be tested for minutes or hours of real time which would be impossible using simulation.

Other components can be written in parallel by other developers, or sequentially by the same developer.

### 3.3 Small Scale Integration

As components are completed they can be integrated. The strong typing, bandwidth allocation, and fixed process creation ensure that components developed by different people will fit together properly. Signal assertions can be written to encode properties (such as signal value or minimum throughput) of the signals, and these can be checked during integration using assertion probes.

If integration fails (components fail to communicate properly), then this is caused by problems between components, rather than within a component (since the component has been verified in isolation, it has static processes, fixed local signals, etc.). The suite of system-wide tools (probes, traces, activity display, etc.) can be used to identify the problem.

### 3.4 Large Scale Integration and Performance Testing

This phase of development can only really be done on the hardware. At this stage all of the FileIO will have been replaced by real components.

It is important to be able to monitor aspects of performance in real-time and this can be done using customised probes which monitor various signals and compare data throughput against predetermined limits. In addition it is possible to monitor the behaviour of the system when processing real-world data, and to inject data by using the microprocessor interface.

The results of the monitoring can be displayed using custom GUI's which the user can develop (an example of a custom GUI is shown in figure 5 in section 5).

### 3.5 Comparison with Traditional Techniques

The power of the overall approach described here is that once a component has been written and tested, it can be assumed to work from that point on. Other parallel systems can behave like this for individual components, but then fail to work during integration, or even worse, during customer use. Possible integration problems include:

- priority inversion (e.g., Mars Pathfinder [7]).
- rogue processes corrupting shared memory
- overflowing message queues
- scheduling failures (e.g., improperly bounded or excessively large critical sections).
- multi-processor bus contention causing non-deterministic communication delays
- incoherency of multi-processor caches

Some or all of these problems will afflict other types of parallel systems, from multi-threaded programs containing just two processes, through to large scale multi-processors. These problems are difficult to track down because they defy logical analysis, they behave non-deterministically and perhaps infrequently, or they disappear when debug code is inserted to find the problem. Even if the cause of the problem is found, it is often difficult to write verification to prevent the bug recurring, because of the need to verify the entire system, not just the component or interface which causes the problem. In the worst case, verification may not be possible because it is not apparent why a fix actually works!

The picoChip solution avoids these types of problem in integration. Individual components, which have been properly verified in isolation, will behave in the same way when integrated into a complete system. The system behaves deterministically, so if problems are found, they can be reproduced, isolated, fixed, and verified. The overall development of systems is more predictable (timescales, etc.), since development isn't held up by strange problems which can't be found until integration.

## 4 picoArray™ Debug and Analysis

### 4.1 Language Features

The language features aid verification and integration through three main features: strong type checking, fixed process creation, and bandwidth allocation.

Strong type checking is used to ensure that whenever data is communicated from one process to another, the data will be interpreted by both producer and consumer in the same way. Types are selected from a library of built-in types, or by the users defining their own types. Types used in communication are limited to 32-bits, which is the maximum size which may be transferred in a single communication over the picoBus™. At the structural level, processes will be defined with ports of specific types, and they will be connected with signals which must match the port types. Within a process, any data which is “put” or “get” from a port must be of the correct type. For processes written in C, this is achieved by synthesising the available types using C encoding rules (e.g., using typedef's, union's, and struct's), and

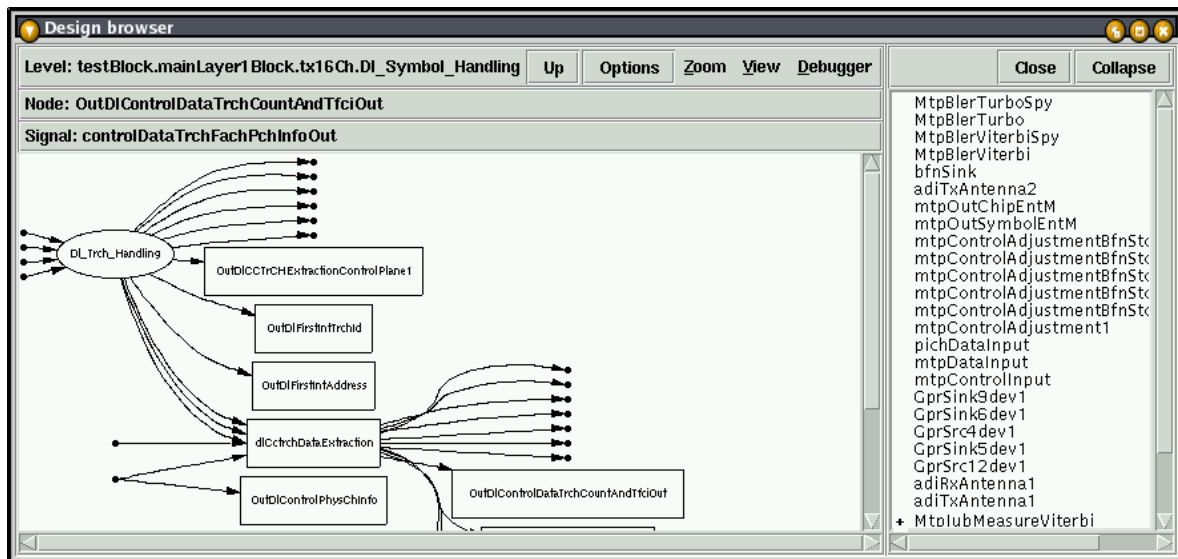


Figure 2: Design browser hierarchical display

hence tying into the C compiler’s type system. Thus, end-to-end communication of data can only occur when all processes and signals agree on the type format. This makes integration of independently developed components easy since any discrepancies in type formats will be detected at compile time, when they are easily fixed.

The structural VHDL used to define a system requires the number of processes, and their interconnections to be fixed at compile time. During compilation, the tools will allocate each process to its own processor, and schedule the signals connecting the processes onto the picoBus™ interconnection fabric. Because of this compile-time scheduling, non-deterministic runtime effects such as process scheduling, or bus contention have been eliminated. This makes it easier to integrate systems. If problems are found, it also makes the reproduction of the problems, their debugging and the verification of their fixes easier.

In addition to specifying fixed signals connecting processes, the signals are also allocated bandwidth. This is achieved using a language notation which allows the frequency of communication over the signal to be specified. Processes requiring high signal bandwidths will use high frequencies (e.g., every 4 cycles), while processes requiring low bandwidth will use low frequencies (e.g., every 1024 cycles).

## 4.2 Design Browser

The design browser is a tool which allows the user’s logical design to be viewed graphically and can be used both during simulation and when executing a design on hardware. The following different graphical views are possible:

- hierarchical,
- flat with a given scope,
- as the strongly connected components (SCC).

The hierarchical view mirrors the structural hierarchy that was created by the user and allows each level of this hierarchy to be explored. An example of this is shown in figure 2.

There are times when the user wishes to see more of a design than is permitted by the hierarchy display and the “flat” display provides this. If displayed from the root of the design the entire design is displayed at once. In addition by displaying from a scope other than

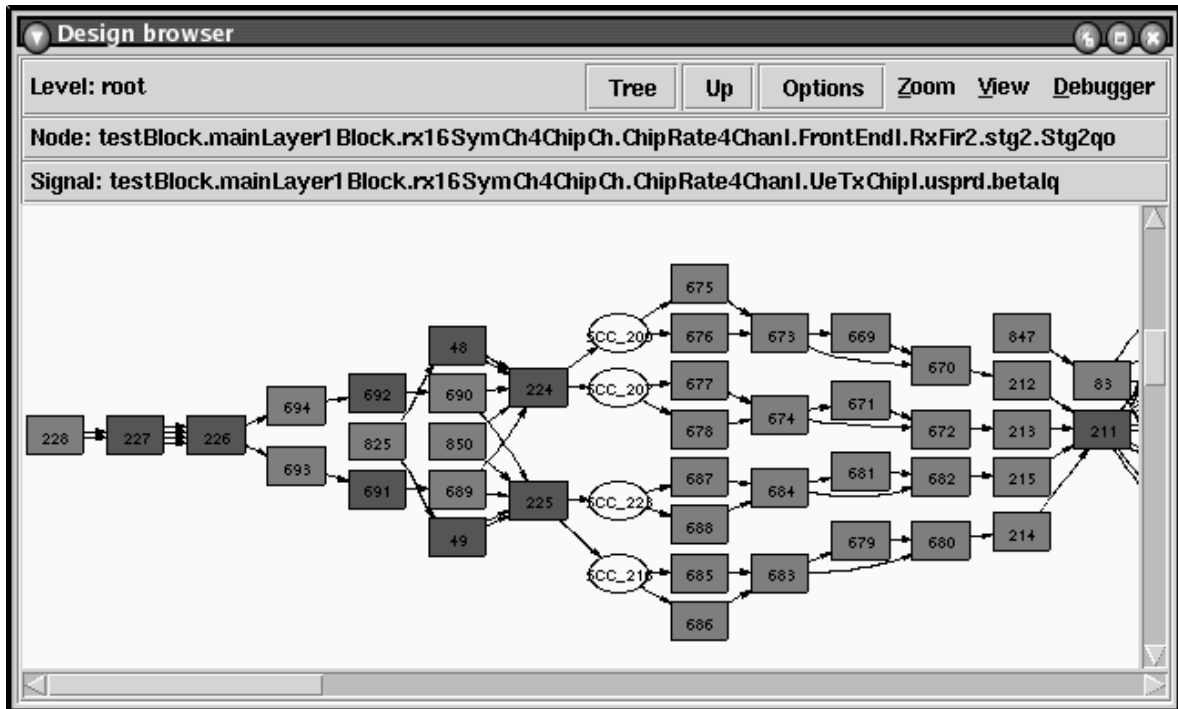


Figure 3: Design browser strongly connected component display

the root subtrees of the design can be viewed. The major difference between this and the hierarchy display is that from a given scope all of the leaf instances are displayed.

The final view comes from thinking of a design as a directed graph and then producing a single level of hierarchy by producing the strongly connected components (SCC). Each of the components can be viewed on their own. An example of this is shown in figure 3. The importance of the SCC view is that from the root level the graph becomes acyclic (directed acyclic graph, a DAG) and therefore this gives advantages when trying to debug a system which has deadlock, livelock or data throughput problems. This separates out the parts of the design that contain feedback from those that are simply pipeline processing.

In addition to these static features the design browser can provide dynamic information about the each instance in a design, for example whether it is processing or waiting for a communications operation. An example of this display is shown in figure 3 (in fact the boxes are coloured, green for processing, red for waiting on communications).

### 4.3 Simulation

The cycle accurate simulation system allows users to build, test and verify their entire design before moving to the hardware. The user is able to extract the state of the system (on a cycle-by-cycle basis) in order to check against the behaviour on hardware. Importantly, the same simulation system was used to provide a “golden reference” during the design and verification of the PC101 and PC102 chips.

The same source-level debugging interface exists on the hardware as on the simulator enabling the user to migrate from one environment to the another without making any changes to their design or their testbenches.



#### 4.4 Scripting

While debugging large parallel systems, operations such as viewing the source code or variable values of individual processes become too low level; this is analogous to debugging a compiled process by inspecting its raw machine code and register values. For large parallel systems it is more convenient to be able to abstract the debugger to provide a higher, system-level interface. Such an interface allows the details of individual processes to be hidden, and replaced by system-specific displays instead. Clearly, it is impossible for picoChip to provide interfaces for every possible system, so instead the debugger can be programmed using Tcl/Tk [8]. This allows the users to create their own system-specific interfaces, built on top of the picoChip debugger. Figure 5 in section 5 shows an example of a WiMax system interface.

#### 4.5 Traffic Analysis

Traffic analysis is used to monitor the state of the communications network. The AEs in a picoArray™ use signals to communicate data (and hence synchronise), and traffic analysis can indicate to the user the states of the communication at any particular time. The relevant data is extracted from selected AEs or from all AEs in a design and either displayed immediately or stored to a file.

For each signal the maximum bandwidth of a signal has to be specified at design time but of particular interest to the user is the actual bandwidth used on a given signal. Using the traffic analysis data can provide information on the statistics of the bandwidth used and can help in the analysis of deadlock and livelock problems.

#### 4.6 FileIO

When testing and debugging it is common to wish to use Unix files in order to inject data into a system or to record intermediate results. This is achieved by providing an AE template which interfaces to the picoBus in the usual way using signals but which is also “connected” to a Unix file. The advantage of this method is that the same user’s code can be used whether the system is running as a simulation or on hardware. The FileIO AE has two different implementations, one for simulation and one for hardware. In a simulation the connection to the file is simple since the simulation simply consists of a piece of compiled C++. In hardware the data memory of the AE is used to buffer the data and when the AE requires it must request that the debugger either empty its memory (for an output FileIO) or fill its memory (for an input FileIO).

#### 4.7 Traces

Traces allow specific types of data, such as register and memory contents, or signal values to be recorded during execution. The trace is stored as a sequence of tuples recording changes in value against the time at which that change occurred. This sequence can be saved to a file, and used by external programs such as gtkwave [9]. The tracing tool is used in a way that is similar to a hardware engineer using an oscilloscope to probe data paths in an electronic circuit. The trace allows a visual representation of the data to be shown with respect to time, which can make certain types of bug readily apparent. Tracing can also be used to perform code profiling, by tracing how the program counter changes over time, and post-processing the information to relate it to the original source code.

While many general purpose processors use special hardware to implement tracing (e.g., ARM Embedded Trace Macrocell [10]) the picoArray™ devices do not. One reason for this

is that traces can generate huge quantities of data (e.g., tracing the program counter for a single processor would generate  $3.2 \times 10^8$  bytes/sec). While the picoArray™ devices have impressive internal communications bandwidth, it would be impossible to transfer this much data off chip without affecting the system being debugged.

Two mechanisms are used to perform tracing. Signals are traced using probes, which are described in section 4.8. The probes mechanism allows signal traces to be performed while running a system at full hardware speed (160Mhz) but the dumping of data to a file means that this speed cannot be sustained. All other types of data (general/special purpose registers, and blocks of memory) are traced using software. The debugger tool repeatedly single-steps the debug system, recording traced values after each step. This can be slow. Typically, the debugged system will be traced off-line, and the results analysed using post-processing tools.

#### 4.8 Probes

Probes are special purpose processes which the debugger inserts into the user's design by utilising unused processors. Probes can be connected to one or more signals, and can non-intrusively monitor all traffic which passes over the signals. They achieve this by using the bus interconnects ability to create 1-to-many connections. For example, suppose two processes in a system were connected by a 1-to-1 signal. If a probe is inserted during debugging to monitor that signal, the debug tools will change the 1-to-1 signal into a 1-to-many signal, with the probe acting as an extra destination. The original processes are unaffected by this change (both in terms of latency and bandwidth), but the probe is now able to monitor all communication over that signal.

Probes are implemented as processes, and so can run at full hardware speed. This enables probes to be used to debug systems in real-time. One use for probes is to allow real-time signal traces to be performed. Other uses include signal assertions, and on-chip analysis.

Signal assertion probes can be used to check that the data passing over a signal conforms to some compile-time specified property. For example, all signals in picoArray™ devices have pre-allocated bandwidth. A signal assertion probe could be attached to a signal to record the bandwidth actually used, thus allowing signals with over-allocated bandwidth to be detected.

Probes can be used to perform on-chip analysis of signal data, rather than having to transport the data off-chip (e.g., using traces), for later analysis. For example, during the development of the picoChip base station, a probe was created which performed Bit-Error Rate (BER) computation on signals. These BER probes could be used to monitor the performance of the base station's Viterbi decoder's in real-time, under different system loads.

#### 4.9 Activity Display

This is related to the trace facility but only looks at the type of activity being undertaken by an AE. This can be running, waiting on a communication or stalled on a memory pipeline fetch. This display allows the history of the activity of a number of AEs to be viewed.

### 5 A Design Example

picoChip is a member of the WiMax (802.16) Forum [11] and is actively working to produce a 802.16 compliant system. The current scope of the work is aimed at producing a system which can be used in either the base station or the consumer premises equipment (CPE) market. As part of this work, picoChip has developed the first part of this system solution - an 802.16 compliant PHY, whose functional decomposition can be seen in figure 4.

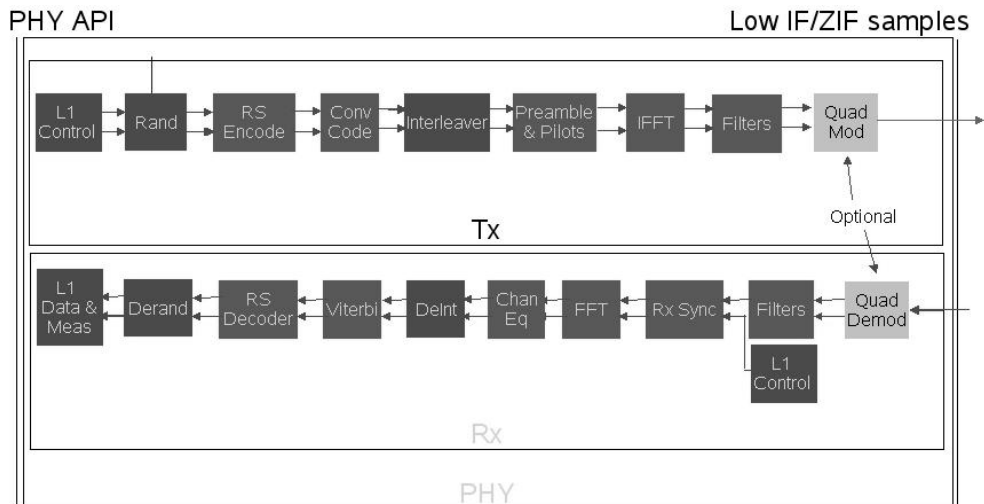


Figure 4: Functional decomposition of an 802.16 PHY

Using the PC101 device the system's team at picoChip have used most of the debugging aids to implement this PHY. This includes developing the individual blocks, executing the implementation on the simulator and creating testbenches in order to verify the correct (compliant) operation before moving onto integration.

It is the final integration of the whole PHY system, for both transmit and receive, that illustrates the key aspects of the debugging environment. The result of this is best shown in figure 5 where the systems group scripted an application specific GUI on top of the primitives provided by the toolkit and indeed the system debugging widgets. There are four areas of interest: three data output probes and one input probe.

The Data Constellation shows data captured by a probe at the output of the Channel Equaliser. This data is extracted in real-time and streamed (using DMA) out of the picoArray via the microprocessor interface.

The Channel Estimation shows the magnitude of preamble sub-carriers (consecutive preambles shown on plot) as captured by a probe at the output of the FFT.

The AWGN (additive white Gaussian noise) has been added to aid checking of the behaviour when there is noise in the channel. This injects data (noise) as input to the quad demodulation block.

The Viterbi BER display shows the Bit Error Rate at the output of the Viterbi, again in real-time, as captured by a probe.

Finally the RSSI (Received Signal Strength Indicator) display shows received signal statistics captured by a probe at the output of the ADI.

## 6 Conclusion

In order to address the target markets in wireless communications picoChip has created a family of code compatible devices (currently PC101 and PC102) using the picoArray™ concept that provides the large computation power required by these applications. To ensure the largest possible computational resource and recognising that system-wide debug is a major problem in multi-processor designs, a rethink of how to apply debugging was necessary. Therefore picoChip has shifted the debug burden from the hardware to the extensible software tools. The tools provide a way of debugging both single processors and more importantly the large multiprocessor systems possible on a picoArray™, or indeed an array of picoArray™ devices. Experience has shown that using these tools picoChip and its cus-

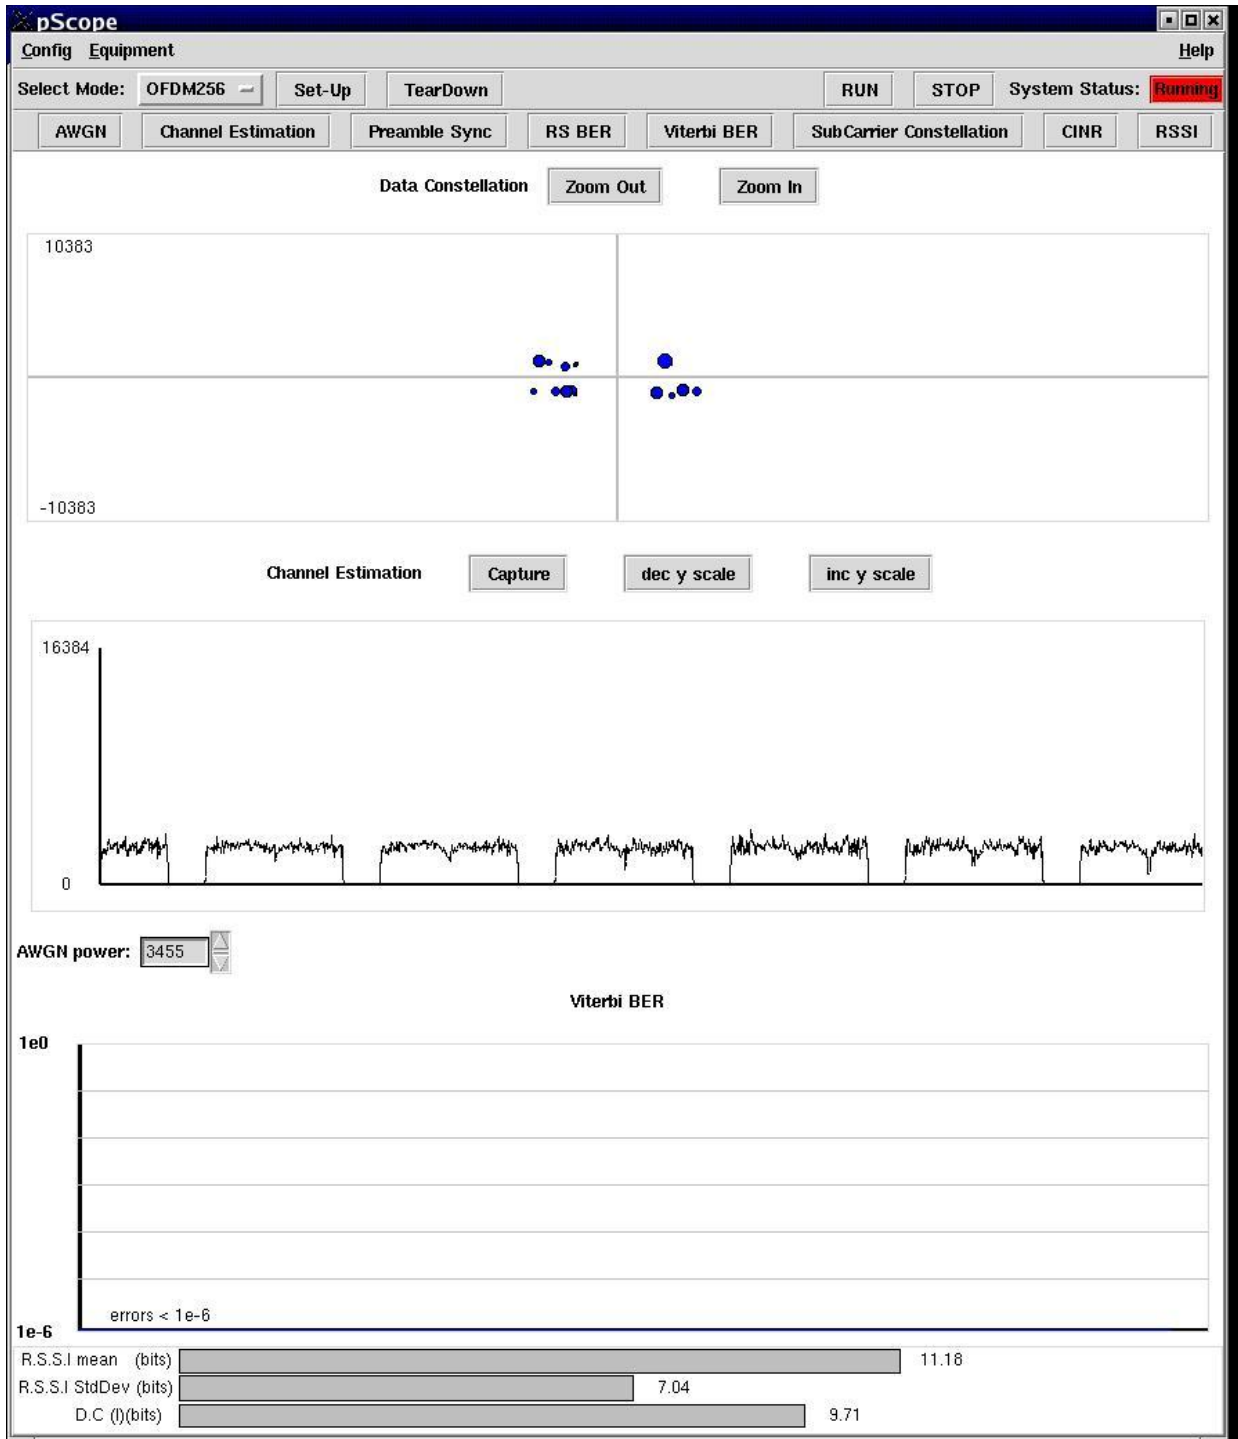


Figure 5: Diagnostics output from 802.16 PHY

tomers have been able to construct large working systems (currently up to 1000 processors) which have been put into our customers' products.

## References

- [1] Andrew Duller, Gajinder Panesar, and Daniel Towner. Parallel Processing — the picoChip way! In J.F. Broenink and G.H. Hilderink, editors, *Communicating Processing Architectures 2003*, pages 125–138, 2003.
- [2] Peter Claydon. A Massively Parallel Array Processor. In *Embedded Processor Forum*, June 2003.
- [3] David R. Butenhof. *Programming with POSIX Threads*. July 1997.
- [4] Peter Pacheco. *Parallel Programming with MPI*. November 1996.
- [5] Sun Microsystems Inc. Threads: Doing Two or More Tasks At Once. <http://java.sun.com/docs/books/tutorial/essential/threads/>.
- [6] Peter Ashenden. *The Designer's Guide to VHDL*. Morgan Kaufmann, ISBN 1-55860-270-4, 1996.
- [7] Mike Jones. What Really Happened on Mars? [http://research.microsoft.com/~mbj/Mars\\_Pathfinder/](http://research.microsoft.com/~mbj/Mars_Pathfinder/), December 1997.
- [8] John K. Ousterhout. *Tcl and the Tk Toolkit*. May 1994.
- [9] The University of Manchester Advanced Processor Technologies Group. GTKWave Electronic Waveform Viewer. <http://www.cs.man.ac.uk/apt/projects/tools/gtkwave/>.
- [10] Embedded Trace Macrocell. <http://www.arm.com/products/solutions/ETM.html>.
- [11] Wimax forum. <http://www.wimaxforum.org/home>.