

Using Application Bisection Bandwidth to Guide Tile Size Selection for the Synchrosalar Tile-Based Architecture

John Oliver¹ and Diana Franklin² and Frederic T. Chong³ and Venkatesh Akella¹

¹ University of California, Davis

² Cal Poly State University, San Luis Obispo

³ University of California, Santa Barbara

Abstract. This paper investigates the impact of proper tile size selection on the power the power consumption for tile-based processors. We refer to this investigation as a tile *granularity* study. This is accomplished by distilling the architectural cost of tiles with different computational widths into a system metric we call the *Granularity Indicator* (GI). The GI is then compared against the bisection bandwidth of algorithms when partitioned across multiple tiles. From this comparison, the tile granularity that best fits a given set of algorithms can be determined, reducing the system power for that set of algorithms. When the GI analysis is applied to the Synchrosalar tile architecture [1], we find that Synchrosalar's already low power consumption can be further reduced by 14% when customized for execution of the 802.11a receiver. In addition, the GI can also be used to evaluate tile size when considering multiple applications simultaneously, providing a convenient platform for hardware-software co-design.

1 Introduction

As power and complexity have become increasingly problematic in modern microprocessors, tile-based architectures have become increasingly attractive (i.e. [2] [3] [4]). In essence, these systems trade architectural complexity for communications, spreading work across a number of sparsely-connected small tiles rather than among richly-connected functional units of a monolithic, wide core.

However, the choice of tile size for tile-based architectures has been largely an ad-hoc, qualitative process. While this may be because of practical reasons (such as availability of cores), this may not yield an efficient design.

In this paper, which is an extended version of our 2006 Computing Frontiers paper [5], we find that in systems where low power operation is critical, proper tile size selection is important. How does an architect find the best tile size for low-power operation? To investigate this we first generate cores with different amounts of computational power. We note that the larger, more richly interconnected tiles have higher average switching capacitance per operation, but also have a larger locality of data available to them. Then, we tile these cores until a fixed amount of total computational parallelism is reached, providing us with a set of tile architectures with different computational *granularity* but with the same amount of total computational power. By mapping the

power efficiency per operation, we then generate a power efficiency curve that we call the Granularity Indicator (GI).

Once we find the GI for a tile architecture with different granularities, we then partition and map different algorithms to the different granularities of the tile architecture and execute them. This process yields the computation cost and the communications cost required for the algorithms to execute across multiple tiles for a tile architecture of differing granularities.

Finally, we can then compare the cost of partitioning an algorithm against the energy efficiency of the tile architecture which is embodied within the GI. If large amounts of communications are exposed by partitioning the algorithm, larger tiles that invest more heavily in connectivity are favored, as they are more apt to hide communications. This is despite the fact that the extra connectivity within larger tiles contribute to higher average switching capacitance. On the other hand, if little communications is exposed when partitioning an algorithm, then smaller, more power efficient tiles are favorable. The result of this comparison is to find the the granularity of tile that has the best power consumption for a given algorithm.

To drive this exploration, we use Synchronscalar [1] architecture as a basis, but other tile-based architectures could be used with a similar methodology. We use the GI framework on the Synchronscalar architecture to investigate how the computation power of the Synchronscalar tiles can be tailored to execute a the 802.11a PHY layer application at low power. We then weigh the cost of this customization against other applications that Synchronscalar may execute. We find that by tailoring the tile granularity to a given application may significantly negatively impact the power consumption of other algorithms, making tile granularity an important decision for low-power tile architectures.

The rest of this paper is structured as follows. First, we develop a set of cores with different amounts of computational parallelism. With these cores, we populate a tile architecture similar to Synchronscalar and generate multiple variants of Synchronscalar with different tile granularities. Using these Synchronscalar-like processor-variants, we can find the GI. Next, we describe the methodology used to partition, map and execute different algorithms on our tile architectures with differing tile granularities. This process yields the computation time required to execute and algorithm and the number of cycles required for inter-tile communication that are required to maintain data coherency across multiple tiles. These cycle results then allows us to compare the communications requirements against the GI. This, we will demonstrate, can tell us which granularity of tile executes a given algorithm at the lowest power. Finally, we use the GI as a guide and we re-design the Synchronscalar architecture for low power 802.11a PHY layer execution. We will show how much power can be saved through the choice of proper tile granularity and also investigate the implications of this optimization on other applications executed on Synchronscalar. We finish this paper with related works and then conclude.

2 Tile Scaling Models

In this section, we develop tiles with different amounts of computational parallelism with which to create a tile-based architecture. While the models presented here could

be developed in a number of different ways, it is important to remember that the central message of the GI arises from non-linear scaling as issue-width grows wider. We argue that this is a valid assumption for any set of tile sizes, since linear scaling would, in essence, reduce a large tile to a collection of small tiles.

To connect the tiles of our tile architecture, we develop models for a bus, statically scheduled mesh, and dynamically scheduled mesh interconnects. These interconnect topologies are intended to be general and cover a wide range of interconnect topologies. Other interconnect networks, such as Raw’s Scalar Operand Network [6], could be employed in a similar study.

2.1 Tile Area

The goal of our tile model is to capture the first-order scaling effects of computational width on area and power. We define the computational width of a tile as the maximum number of arithmetic operations that can be completed per clock cycle, where the operands are in the local register file. The smallest tile we consider in this study can compute a single operation in every cycle, while the largest tile we consider can compute 32 operations in parallel every cycle. We assume a VLIW-based architecture, which can be efficiently scheduled for data-parallel applications like media applications. The register file of our model is assumed to provide one write and two read ports for each operation.

We first developed a tile based on the Blackfin Digital Signal Processor (DSP) [7], which can be viewed as having a computational width of two. In order to get a power and area estimate for this processor, we modeled the control logic of this processor in VHDL and synthesized it using the Synopsys Design compiler. The data-path units, i.e. multipliers, register file and memory, were estimated using published numbers [8–10]. Parameters used to develop a model of the Blackfin DSP are shown in Table 1.

Table 1. Technology Parameters

Parameter	Value	Source
Technology	130 nm	
Minimum Voltage	0.7V	Blackfin DSP [7]
Maximum Voltage	1.65 V	Estimated [11]
Threshold Voltage	0.332 V	[11]
Temperature	40 C	Assumed
Oxide Thickness	3.3 nm	[11]
Dielectric Strength of Oxide	5e6 V/cm	[11]
Max Frequency	600 MHz	SPICE using [11]
Tile Power	0.1mW/MHz	See estimate above
Tile Size	1.82 mm^2	Synthesis
Wire Cap.	387 fF/ μm	Semi global [12]
Wire pitch	16 λ	Semi global wiring [12]

Using the width-2 Blackfin DSP as a basis, we extrapolate the tile area for tiles with a computational width of one, four, eight, sixteen and thirty-two. We assume that the

area of control logic scales linearly with computational width as well as the area contributions of the ALU, shifter, accumulator and multiplier. Memory capacity is assumed to grow linearly with computational width at 32 KB of instruction and 32 KB of data memory per computational width.

For the register file, we assume that the number of ports in the register file, as well as the capacity, grows linearly with the computational width. This produces a quadratic increase in both power and area in the register file. Finally, the on-chip wiring/data-forwarding paths are also assumed to grow quadratically, in a similar manner as the register file.

Figure 1 shows the area results of our tile model. We hold the total computation width constant at 32 computational widths, so when we halve the computational width of a tile, we double the number of tiles we are using. The left most bar in Fig. 1 shows the area breakdown for a single tile with a computational width of 32. The next bar to the right shows a tile model with two tiles with width 16. The column furthest on the right shows the area of 32 tile each with a single computational width. The single large tile with a width of 32 has a 93% area growth over the array of 32 tiles with computational widths of one. Note that if processor area is a design constraint, this will need to be weighed in conjunction with any power saving we present in this paper. However, for this study, we concentrate only on saving power.




Fig. 1. Tile area scaling for one 32-wide tile (1:32) to 32 1-wide tiles (32:1).

2.2 Tile Power

The tile power of our tile model is composed of two portions, the active power and the leakage power. To find the active power, we use power numbers based upon synthesis of the Blackfin core, as described in the previous section. This yielded a power estimate of 0.1 mA/MHz at 1 V, on average.

For the other granularity of tiles the average current for each of tiles is assumed to be proportional to the area relative to our Blackfin core. This is a decent approximation if two conditions are met. First, the activity factor of the tiles must be constant. Since the partitioning of data parallel multimedia applications used in this study are done in a load-balanced manner, this should hold approximately true. Second, for those

micro-architectural structures that have non-linear area growth, their power consumptions must track the growth in area. This is true for the register file to a first approximation, as register files have been shown to have active power consumption that is linearly proportional to area [10].

For our leakage model, we assume that leakage power is proportional to the number of transistors. Using an average of 830 pA of leakage per transistor [13], we approximate that the Blackfin DSP leaks 1.5mA. This provides a range where the smallest, single-width tile leaks 0.74 mA of current, and the largest, width-32 tile leaks 23.68 mA of current.

Having established the assumptions for our tile power scaling, we find that this provides a range of currents consumed for different sized tiles. A single computational width tile uses 0.05 mA/MHz on average, while the largest 32-width tile consumes 4.87 mA/MHz on average. For a total of 32 computation ways, this yields a tile architecture that has current requirements as shown in Fig. 2.




Fig. 2. Current required for our tile architecture model with 32 computational widths of differing granularities.

2.3 Tile Power Model Correlation

Although the GI metric and analysis methodology can be applied to a wide array of tile power scalings, in order to demonstrate the usefulness of the GI, our tile model needs to reflect the scaling trends that real processors will observe. In order to see if our tile power model scales as industrial processors do, we have plotted published power results from similar processors from industry. We expect our power models to lie below the curve of the realistic processors for two reasons. First, we model only the core components, not the I/O devices and special purpose circuits. In addition, often commercial projects scale not only the width, but also the functionality, adding specialized units and other functions while we are only looking at computational cores. Figure 3 shows commercial processors, normalized for process technology. Next to Fig. 3, we show Table 2 which contains the references for the processors used in Fig. 3. We can see that, as expected, our scaling model shows a similar trend but has lower absolute power than the published results.




Fig. 3. Current required for our tile architecture model with 32 computational widths of differing granularities.

Table 2. Published numbers for similar VLIW-based processors

Plot #	Processor	Citation
1	Analog Devices ADSP-2191	[14]
2	TI TMS320C2810	[15]
3	NEC SPXK5	[16]
4	Hitachi SH-Mobile3	[17]
5	Infineon Tricore 2	[18]
6	Analog Devices TS-101	[19]
7	Transmeta Crusoe TM5700	[20]
8	Transmeta Efficeon TM8820	[21]
9	TI TMS320C62x	[22]
10	TI TMS320C64x	[23]
11	TI TMS320DM643	[24]

2.4 Inter-Tile Interconnect

To properly account for the power due to inter-tile communication, we need two values - the delay caused by the interconnect (and thus the idle cycles of the tiles) and the power required by the interconnect to perform communication. In this study, we evaluate both a single bus and a generalized mesh interconnect topologies. We also assume that data is communicated between tiles using explicit message passing. A similar study could be done with a shared memory system with a hardware-enforced coherence protocol.

Interconnect Delay In order to calculate communications delay, we must know how much communication occurs and how long each transmission takes. Details on how we find the amount of communication required by an algorithm are described later in Sect. 4. For a shared bus, the delay for each communication is a single cycle, regardless of the source and destination. Since the distance of the bus is small and operational frequencies are limited in the Synchronscalar architecture, single cycle communications is possible using a bus. The delay of a mesh is a function of the contention on the mesh. This requires a traffic simulator to accurately find contention on the mesh. Mesh simulations were completed using the FlexSim mesh simulator from USC [25]. FlexSim was configured in a 2-D space for up to 32 switches, where each switch is attached to an end-node with one injection channel to the switch. FlexSim was modified in two ways. First, the default latencies were reduced to allow low-overhead flit-level routing as expected for an on-chip network. Also, an optional mode was introduced in which only the link overhead was counted, and the routing overhead was discounted, in order to simulate the delay for a statically-scheduled mesh. This allows us to more closely emulate the statically scheduled nature of the Synchronscalar inter-tile interconnect.

Interconnect Power For our inter-tile interconnect power model, we employ power costs as abstracted from the Orion interconnect power model [26] from Princeton University. We find that our wires are using in the neighborhood of 10 pJ/bit for a 10mm trace, similar to Stanford’s Smart Memories [27].

As observed in previous studies by the RAW project [28] and the Synchronscalar project [1], interconnect switching power can be a small portion of the overall power consumption. The reason for this is two-fold. First, the number and size of tiles used in this study are relatively small, thus not requiring an abundance of interconnect resources. Second, the frequencies of operation of the tiles in this study is relatively low compared to high-speed processors, resulting in low frequency communications as well.

3 The Granularity Indicator (GI)

Now that the tile area and power model used in this study has been introduced, we will encapsulate the power of different granularities of tiles into a metric we call the Granularity Indicator (GI).

The GI expresses the *architectural power characteristics* of a tile architecture that is comprised of tiles with different computational granularities. At its simplest form, the GI is a measure of the relative energy efficiency per operation of different granularities of tile architectures, similar to Fig. 2 except with a simple added transformation.




Fig. 4. The GI for tile our model. A single 32-width tile is shown on the left and thirty-two 1-width tiles is shown on the right. Smaller tiles have less average switching capacitance, which can be re-invested into communications.

The additional energy saved on every operation by a finer granularity tile architecture is reinvested into a communications budget. So, for every pJ of energy a smaller tile saves in energy consumption versus a larger tile, it can re-invest that energy into communications. This has the effect of changing the vertical axis of Fig. 2 from power consumption to allowable communications overhead while maintaining iso-power consumption.

For our tile model, with an power consumption curve as shown in Fig. 2, this transformation creates the GI which is shown in Fig. 4. On the left of Fig. 4, we see a single large tile with 32 computational widths. As we move to the right on Fig. 4, we double the number of tiles but halve the widths. By moving to finer-grain tiles, we know that the average energy consumption per operation is reduced, as shown in Fig. 2. However, this is shown as allowable communications given iso-power consumption in Fig. 4. This change in axis is convenient for comparing the energy efficiency of a tile architecture with a given granularity for a given communications requirements for a given algorithm because the communications overhead allowed can be easily matched against the communications requirements of a partitioned algorithm.

Alternatively, to find the system power of an architecture executing a given algorithm, extensive simulation is typically involved. This transformation to the GI allows us to decouple the architectural contributions to power consumption from the algorithm’s demands for computation and communications cycles. This decoupling allows an architect to make architectural-based decisions to minimize power while quickly evaluating the effectiveness of those architectural decisions for a given set of algorithms. In Sect. 3.1 we look at the impact of voltage scaling on the GI. In Sect.3.2 we describe how a mode that puts tiles into a low-power mode when completing inter-tile communications can affect the GI. In Sect. 5, it will be clear how these shifts can lead to understanding of the effectiveness of certain architectural features for a given algorithm or application.

3.1 Tile Voltage-Frequency Scaling

As we can see from the GI in Fig. 4, a tile architecture comprised of smaller tiles has a higher allowable communications overhead than a tile architecture made of larger tiles and still have the same power consumption. However, to support this communications overhead, additional cycles are required which will result in a higher operational frequency for a given throughput. Additionally, this higher frequency requires a higher op-

erational voltage, which also will increase the power consumption of a tile architecture comprised of smaller tiles. Therefore, in the presence of voltage scaling, a tile architecture with many small, higher frequency tiles will consume relatively more power than a tile architecture with fewer, lower frequency tiles.




Fig. 5. Finer grain tiles require more cycles to execute a given algorithm, because of added communications costs. Added cycles require higher frequency of operation, thus higher supply voltages.

Figure 5 shows four GI curves, corresponding to four different base frequencies. The base frequency is defined as the frequency of operation of a single tile with a computational width of 32. We see that the GI is shifted down-wards at higher frequencies. So the impact of voltage scaling is to reduce the amounts of allowable communications for a finer grain tile architecture for the same amount of power consumption as a coarser grain tile architecture.

3.2 Low Power Idle Tiles for Low Power Communication

The GI can also show how using an idle communication mode in the tiles impacts the power consumption. The Blackfin DSP requires a total of two cycles to enter two cycles to exit from idle mode. While in idle mode, the core consumes approximately a fifth of the active power consumption. For communications that cannot be overlapped with computation, this mode can be used to reduce the overall system power.

Figure 6 shows the impact on the GI when using tiles with this idle mode for non-overlappable communications. Intuitively, communication with the addition of this mode now costs relatively less. The result is that fine-grain tiles should become more attractive since fine-grain tiles require more communication than coarse grain tiles for a fixed amount of aggregate parallelism. Likewise, we would expect the power consumption of an application mapped onto fine grain tiles to decrease. This effect is shown on the GI in Fig. 6. We see that this can make a dramatic difference in the communication supported by smaller tiles. Not surprisingly, as more tiles (requiring more communication) are used, more power is saved by the implementation of the idle mode.

Now that we have introduced the GI and presented a pair of features that may impact the GI, we will now talk about how we find the communications overhead required by partitioned algorithms.




Fig. 6. An idle mode used to limit the energy consumption of tiles when idle. Since finer grain tiles expose more communications, features that reduce the power cost of communications favor finer grain tiles. The result is that the GI shifts up-wards.

4 Algorithm and Applications Partitioning Methodology

The goal of our algorithm partitioning and mapping analysis is to find out, for each granularity of tile architecture, the amount of communication and computation that needs to occur for completion of that algorithm and to minimize this communication requirement.

We use several typical media algorithms in this study. All of the algorithms evaluated in this paper are statically schedulable media algorithms. We assume that these algorithms are partitioned in a latency sensitive manner, so algorithms such as FFT and FIR can be used in applications such as the IEEE 802.11a wireless protocol.

FIR filter evaluated is a 64 TAP, 16 sample block FIR filter.

FFT used a radix-2 64 point transform.

Viterbi decoder implemented is a constraint length = 7, 208b block 1/2 rate decoder.

LDPC decoder is a block error decoder. From a computer architect's standpoint, this can be viewed as a huge message-passing algorithm that requires over 61000 messages to be passed between the nodes of a bipartite graph.

Software Radio is a benchmark by the MIT Raw project. Each sample is read in, filtered, demodulated and equalized with the previous inputs.

802.11a 54 Mbps wireless receiver. The four major components in the 802.11a receiver are the FFT, demodulation, de-Interleaving and the Viterbi decoder. The FFT is a 64-point FFT, demodulator and de-Interleaver is implemented by a look-up table and filtering. The Viterbi Decoder is a K=7 one half rate decoder.

MPEG-4 encoder. We implement Motion Estimation (to remove temporal redundancy between successive frames), DCT and Quantization, which constitute about 90% of the computation of the video encoder [29].

The algorithms we chose to evaluate are those that can be executed on the Synchronous tile architecture, namely static media-based applications. Due to the static nature of these algorithms, we adopted a graph-oriented approach using the best-known

graph partitioning algorithms to obtain the best parallelization possible for each granularity. For dynamic workloads, other partitioning methods may be used in conjunction with the GI to find the optimal power consumption of those workloads.

The first step is to express the algorithms as data flow graphs (DFGs). Next, an algorithm is then partitioned onto multiple tiles. To partition and map the DFGs, we iteratively employ Chaco [30], which is a graph partitioning tool that is used in the scientific computing community for high-performance multiprocessors. In particular, Chaco uses recursive spectral bisection (which performs minimum cuts through eigenvalues of an adjacency matrix) with a Kernighan-Lin heuristic to improve the partition resolution. The result is load balanced partitions with minimal N-section bandwidth.

To find the execution time of the algorithms, we can first use the Blackfin simulator to find the execution time of the computational nodes in the graph. Then we use the FlexSim cycle accurate network simulator [25] to simulate the communications cycles required. This process is repeated for each tile granularity and each algorithm, yielding the number of computational and communications cycles required for each algorithm on each granularity of tile architecture.

5 Partitioning Results and Granularity Analysis

We now show the partitioning results of several algorithms and compare these results against the GI. This will allow us to see which algorithms execute most efficiently on a given granularity of tile architecture.

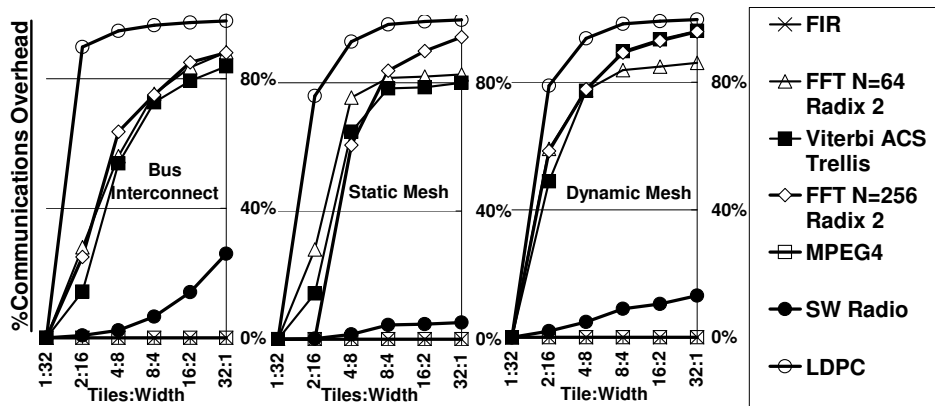


Fig. 7. The communication overhead of our algorithms on three interconnection networks. The left chart shows results for a bus interconnect, the middle for a statically scheduled mesh, and the right for a dynamically scheduled mesh.

Figure 7 shows the amount of inter-tile communication overhead of the different algorithms on three different interconnects - bus, statically routed mesh, and dynamically routed mesh. We assume that the interconnect wire-widths is 32b for each of the topologies. We can see in Fig. 7 that LDPC requires the most communication, and

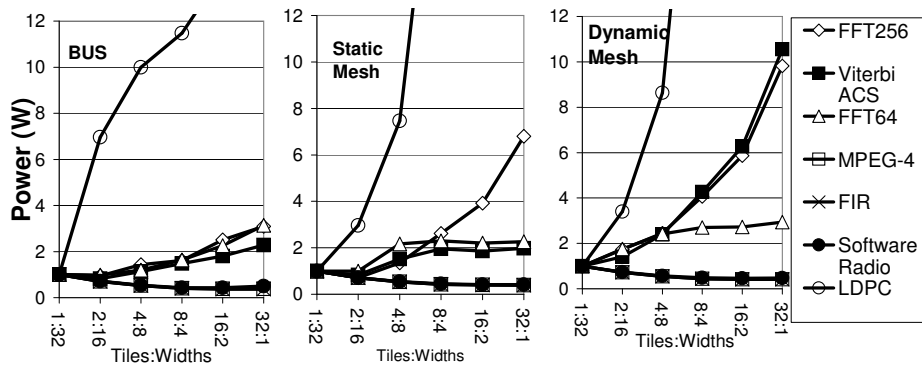


Fig. 8. The relative power consumptions of our algorithms on three interconnection networks. Note that these are relative to the difference between the communications requirements of an algorithm and the GI.

both trellis-based algorithms, the FFT and Viterbi ACS, have high degrees of communication. For an FIR block filter, our partitioning was able to hide most of the required communications. Likewise, MPEG4 encoding and Software Radio do not expose large amounts of communications when partitions cross multiple tiles.

Now, let's compare the communications requirements of these algorithms against the GI. The dotted line in Fig. 7 is the GI for our tile scaling model for the Synchronscalar architecture. Remember, the GI curve shows the maximum amount communications while maintaining iso-power consumption for tile architectures with different tile granularities. Therefore, algorithms that have inter-tile communications requirements above the GI for a given tile granularity will execute more efficiently a coarser-grain tile architecture. If the exposed communications curve is above the GI at all granularities, then a single large tile is the most efficient choice. Conversely, algorithms with communications requirement below the GI will execute most efficiently at the granularity that has the largest distance between the GI curve and the exposed communications curve. It is in this way that the GI exposes relative power consumptions for different granularities of tiles for a given algorithm.

For instance, in Fig. 7, we can see that on a static mesh interconnect, the 64 point FFT (as marked by white triangles) is above the dotted GI line at all tile granularities except for a single large tile. This indicates that the amount of communication exposed by partitioning the FFT is greater than what is allowed by the GI to maintain iso-power execution. Therefore, a single large tile is the most efficient granularity for a 64 point FFT that requires 32 widths of total performance. Likewise, we can see that the Viterbi ACS trellis (as marked by black squares), dips below the GI when executed on a static mesh for two tiles with a computational width of 16. Since the Viterbi ACS exposed communications curve is above the GI at all other granularities, the Viterbi ACS trellis will execute at lowest power on two tiles with a computational width of 16.

To show that the relative location of the communications curve compared with the GI is a good indicator of minimal power consumption, we also show in Fig. 8 the power consumption of these algorithms on Synchronscalar-based tile architecture. In Fig. 8, for a static mesh, we can see that FFT (again marked by white triangles) has a higher

power consumption for all partitionings, so a single large tile is best. For Viterbi ACS, the power consumption is lowest for the 2:16 case for the statically scheduled mesh, just as indicated by the GI analysis.

Finally, to summarize, Fig. 9, we show the power consumption of six different applications, the Viterbi ACS, MPEG4, Software Radio, 64 point FFT, a 64 TAP FIR filter, and the 802.11a application at 6 Mbps. The granularity of tile with the lowest power consumption is circled for each application or algorithm.




Fig. 9. Power consumption of different algorithms for different tile granularities.

5.1 Using the GI

Thus far, we have developed the GI to describe the best-fit granularity of tiles for a given application. However, the GI can also be used to direct architectural design.

Suppose an architect is interested in building a media processor to run FFTs. Furthermore, the architect has flexibility to choose what size of tile to use as well as the width of the mesh-based interconnect, but is constrained to a total of 32 computational widths. The GI can be used to guide the architect to these decisions.

Figure 10 shows the communication overheads for a 64 point FFT, mapped with the base-line GI. From this figure, we can see that the FFT requires more communication than a 64b mesh can support at any granularity, except for a single large tile. Therefore, for a tile architecture with a 64b mesh, running on a single tile is the best option. However, as we increase the bandwidth of the mesh, the inter-tile communication overhead

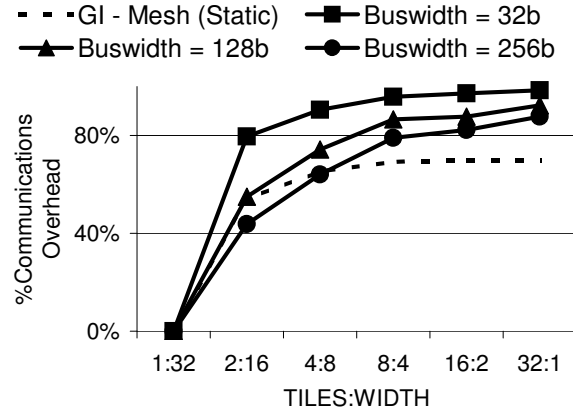


Fig. 10. The GI and the communication requirements for a 64 point FFT for different mesh bandwidths.

is reduced. This has the effect of making large tiles that have high amounts of local on-tile interconnect relatively less powerful than smaller tiles. Indeed, for a 256b mesh, we see that the FFT's communication overhead curve has dipped below the GI for 2:16 and 4:8 points. This indicates the large 32-wide tile is no long the most power efficient. Furthermore, by utilizing the GI for the FFT with our tile model, we can make trade-offs between tile granularity and inter-tile interconnect. These trade-offs are shown in Fig. 11. For instance, we can see that 32 width-1 tiles using a 128b mesh has a lower GI number than eight width-4 tiles using a 64b mesh. The GI Number gives the architect the ability to weigh the added (predominantly) area cost of the larger interconnect with the power saved by using a larger mesh.

Likewise, we can see the power consumption of a few end-to-end applications. Figure 12 shows the power consumption of the 802.11a application on different granularities of tiles with different bandwidth interconnects. As expected, for a communications intensive application like 802.11a (and thus an application who's bisection bandwidth is above the GI), adding more interconnect bandwidth significantly improves the power consumption of the system. From Fig. 12, we see that moving from 64b to 128b interconnects, marked by squares and triangles respectively, decreases the power consumption greatly. The impact is much larger for the coarse granularity tiles because communication stall cycles are much more expensive for coarse-grain tiles than fine-grain tiles. Figure 13 shows a similar result for the software radio application.

6 Application of the GI to Improve Synchronoscalar

Using the GI, we can revisit the design of the Synchronoscalar architecture. Synchronoscalar was designed as a system based on 2-wide Blackfin tiles. We will attempt to use the GI

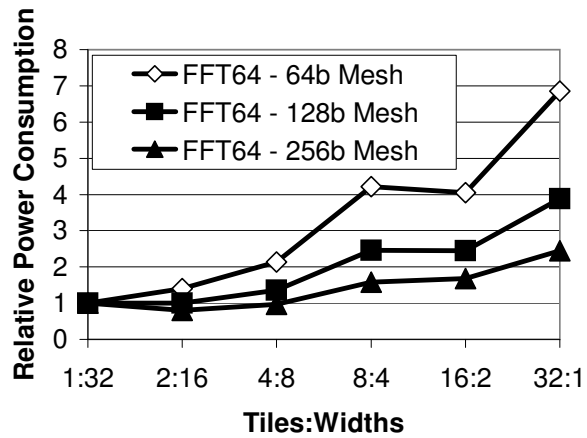


Fig. 11. Relative power for FFT plotted for different mesh bandwidths. We can see that the bandwidth of the inter-tile interconnect impacts the best granularity of tile for low-power execution. For a 256b Mesh two 16-width tiles is most efficient, while for a 64b mesh, a single 32-width tile is most efficient.

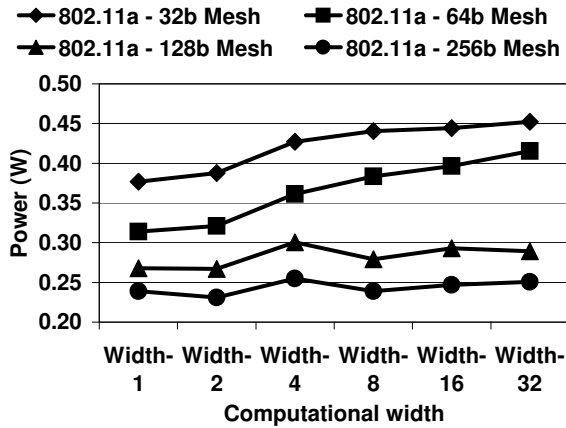


Fig. 12. The power consumption of the 802.11a application on different granularity of tiles with different bandwidth interconnects.

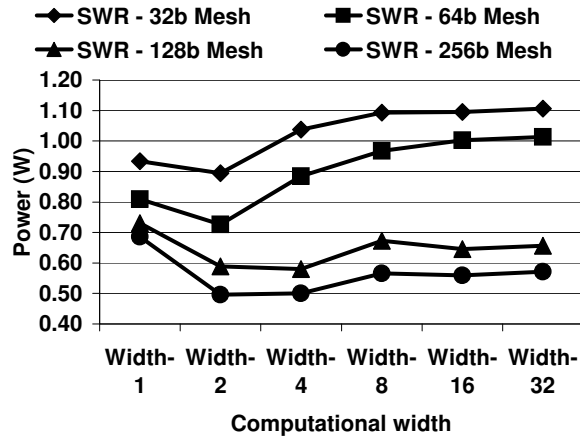


Fig. 13. The power consumption of the software radio application on different granularity of tiles with different bandwidth interconnects.

and customize the tile granularity for efficient execution of 802.11a PHY layer baseband processing.

We will assume the same tile scaling model as presented in Sect. 2. First we need to compute the GI for our tile model. We will also assume that this version of Synchrosalar uses a generalized, statically scheduled mesh as an inter-tile interconnect. The GI for this model is shown in Fig. 14 as a dotted line.

Next, we need to find the communications overhead of 802.11a when it is partitioned across 2, 4, 8, 16 and 32 tiles. This is shown in Fig. 14 as a solid black line.

Now, to find the most efficient granularity, all we need to do is find the place where the communications exposed by partitioning 802.11a is lowest relative to the GI. From Fig. 14, we can clearly see that this occurs with two tiles with a computation width of 16. Again, for validation, the power consumption of 802.11a on differing granularities of Synchrosalar has also been plotted in Fig. 14, as marked by triangles. Indeed, we can see that two 16-width tiles is the lowest power consuming granularity, saving Synchrosalar 14% power over Synchrosalar’s already very low power consumption.

Before moving on, in Sect. 3.1 and 3.2 we discussed how architectural features can shift the GI up or down. We can now see how this is useful information. A downward shift in the GI would likely mean that for the Viterbi decoder on a static mesh, two 16-width tiles would no longer be the most efficient operating point. This is because if the GI shifted down-wards (perhaps by implementing dynamic voltage-frequency scaling on the tiles), the Viterbi decoder’s communications requirements would likely be above the GI curve. So, if the GI shifted down-wards, this indicates that a single large tile would be the most efficient for executing the Viterbi decoder. A similar result can be seen here in the case of 802.11a. Conversely, if tile idle modes were implemented, the GI would shift up-wards. This would perhaps allow some more algorithms to execute more power efficiently on finer grain tiles.

Now that we have found the tile granularity that most efficiently executes for 802.11a, lets investigate the impact of this on the other applications that Synchroscalar supports. In Fig. 15, the power consumptions of four different applications are shown, both for the original width-2 tile Synchroscalar array and the width-16 tile Synchroscalar array. The changes in power consumption are shown numerically on top of each pair of bars. We see that 802.11a saves about 14% power over the orginal Synchroscalar array, but this comes at a cost of a 65% increase in power consumption for MPEG4. The architect can then easily find the best trade-off of tile granularity and power consumption for all the applications of interest by using the GI as a guide. This makes the GI a useful hardware/software co-design tool.




Fig. 14. The dotted line is the GI and the solid Line is the communications overhead exposed when partitioning the 802.11a signal chain on up to 32 tiles. The Power Consumption of 802.11a tracks the relative distance between the GI and communications overhead.

7 Related Work

Our work attempts to build intuitive understanding of a design space occupied many diverse projects. The MIT SCALE project [31] is developing a tile-based power efficient architecture based on their Vector-Thread paradigm. In their prototype SCALE processor, they are able to develop a simple micro-architecture that attains high performance and low power execution by avoiding complex control structures and utilizing spatial locality. The EnyAC group at Carnegie Mellon [32] is investigating globally asynchronous, locally synchronous designs to allow for dynamic voltage and frequency scaling for low power consumption.

On the processor-power efficiency front, Zyuban [33] has developed an architectural based power-performance efficiency metric for a single microprocessor which allows efficiency to be evaluated during the development of the ISA. In a complementary




Fig. 15. Four different applications are plotted for the original width-2 Synchroscale array and the width-16 Synchroscale array. Added power consumption percentages are shown above each set of bars.

study, Hartstein and Puzak [34] develop a power efficiency metric and investigate the power efficiency of deep pipelines on a processor. While these studies are concerned with the power efficiency of a single tile, our study extends the study of power efficient processors to multiple-processors on a single chip. In the paper, Custom Fit Processor [35], a VLIW tile model is developed and performance is weighed against area cost, but not against power. A study similar to ours for energy efficient interconnects [36] has been published by Heo and Asanovic.

Finally, it is because of the many different tile based architectures that are being researched that this study was developed. The RAW project [2] [37] uses MIPS-based cores as tiles and shows performance scalability through their robust, three-level inter tile communication structure. Also, in a similar effort is the Smart Memories project *smart-memories*. Smart Memories uses finer-grain tiles than the RAW processor. The TRIPS architecture [3] also attacks wire-scalability by utilizing multiple cores. Additionally, TRIPS is a malleable architecture that can adapt to different types of workloads to gain performance, yet maintain performance for general purpose workloads. One study that looks at a heterogeneous tile structure was done at Technion [38] and allows the core with the best power efficiency to execute.

8 Conclusions

The Granularity Indicator (GI) provides a novel way to encapsulate power scaling factors when trying to meet performance targets with parallelism. The GI can be used to discover which algorithms can be executed in a power efficient manner on small or large tiles. Additionally, through the use of the GI and knowledge of communication overheads from algorithms, tile architectures can be optimized for granularity of targeted application mixes.

We have presented the GI and used it to show many different forms of analysis. We have explored how base frequency and idle-modes affect the power-performance

scaling and how applications behave with different tile widths. Finally, we used the GI to revisit tile granularity in Synchronscalar. We found that the use of the Blackfin DSP as our tile, was non-optimal in terms of power for our centerpiece application, the 802.11a receiver.

References

1. Oliver, J., Rao, R., Sultatna, P., Crandall, J., Czernikowski, E., Jones, L.W., Franklin, D., Akella, V., Chong, F.T.: Synchronscalar: A multiple clock domain, power-aware, tile-based embedded processor. In: 31th Annual International Symposium on Computer Architecture (31th ISCA-2004) Computer Architecture News, ACM SIGARCH / IEEE (2004)
2. Taylor, M.B., et al.: The Raw microprocessor: A computational fabric for software circuits and general-purpose programs. *IEEE Micro* **22**(2) (2002) 25–35
3. Sankaralingam, K., Nagarajan, R., Liu, H., Kim, C., Huh, J., Ranganathan, N., Burger, D., Keckler, S.W., McDonald, R.G., Moore, C.R.: Trips: A polymorphous architecture for exploiting ilp, tlp, and dlp. *ACM Trans. Archit. Code Optim.* **1**(1) (2004) 62–93
4. Mai, K., Paaske, T., Jayasena, N., Ho, R., Dally, W.J., Horowitz, M.: Smart memories: A modular reconfigurable architecture. In: 27th Annual International Symposium on Computer Architecture (27th ISCA-2000) Computer Architecture News, Vancouver, British Columbia, Canada, ACM SIGARCH / IEEE (2000) Stanford Univeristy.
5. Oliver, J., Rao, R., Brown, M., Mankin, J., Franklin, D., Chong, F.T., Akella, V.: Tile size selection for low-power tile-based architectures. In: CF '06: ACM International Conference on Computing Frontiers. (2006)
6. Taylor, M.B., Lee, W., Miller, J., Wentzlaff, D., Bratt, I., Greenwald, B., Hoffmann, H., Johnson, P., Kim, J., Psota, J., Saraf, A., Shnidman, N., Strumpfen, V., Frank, M., Amarasinghe, S.P., Agarwal, A.: Evaluation of the raw microprocessor: An exposed-wire-delay architecture for ilp and streams. In: ISCA. (2004) 2–13
7. Kolagotla, R., Fridman, J., Aldrich, B., Hoffman, M., Anderson, W., Allen, M., Witt, D., Dunton, R., Booth, L.: High Performance Dual-MAC DSP Architecture. *IEEE Signal Processing Magazine* (2002)
8. Gupta, S., Keckler, S., Burger, D.: Technology independent area and delay estimates for microprocessor building blocks. In: Technical Report TR2000-05, Department of Computer Science, University of Texas. (2000)
9. Wolfe, A., Fritts, J., Dutta, S., Fernandes, E.S.T.: Datapath design for a vliw video signal processor. In: HPCA. (1997) 24–
10. Rixner, S., Dally, W., Khailany, B., Mattson, P., Kapasi, U., Owens, J.: Register organization for media processing. In: International Symposium on High Performance Computer Architecture (HPCA), Toulouse, France (2000)
11. Hu, C.: Berkeley predictive technology model (2004)
12. Ho, R., Mai, K., Horowitz, M.: The future of wires. In: Proceedings of the IEEE. Volume 89. (2001) 490–504
13. Thompson, S., Alavi, M., Hussein, M., Jacob, P., Kenyon, C., Moon, P., Prince, M., Sivakumar, S., Tyagi, S., Bohr, M.: 130nm logic technology featuring 60nm transistors, low-k dielectrics, and cu interconnects. *Intel Technology Journal* **6**(2) (2002) 5–13
14. AnalogDevices: ADSP-2191 Processor Data Sheet. (2002)
15. TexasIntruments: TMS320C28x Processor Manual. (2001)
16. T. Kumura, M. Ikekawa, M.Y., Kuroda, I.: VLIW DSP for Mobile Applications. *IEEE Signal Processing Magazine* (2002)

17. Mizuno, H., Irie, N., Uchiyama, K., Yanagisawa, Y., Yoshioka, S., Kawasaki, I., Hattori, T.: SH-Mobile3: Application Processor for 3G Cellular Phones on a Low-Power SoC Design Platform. *Hot Chips 16* (2004)
18. Norden, E., Leteinturier, P., Barrenscheen, J., Scheibert, K., Hellwig, F.: A Fast Powertrain Microcontroller. (2004)
19. AnalogDevices: TS-101 Data Sheet. (2002)
20. Transmeta: Transmeta Crusoe TM5700/5900 Processors. (2003)
21. Transmeta: Transmeta Crusoe TM8300/8600 Processors. (2004)
22. TexasInstruments: TMS320C62x Processor Manual. (2001)
23. Agarwala, S., Anderson, T., Hill, A., Ales, M.D., Damodaran, R., Wiley, P., Mullinnix, S., Leach, J., A., L., Gill, M., Rajagopal, A., Chachad, A., Agarwala, M., Apostol, J., Krishnan, M., Bui, D., An, Q., Nagaraj, N.S., Wolf, T., Elappuparackal, T.T.: A 600MHz VLIW DSP. (2002)
24. TexasInstruments: TMS320DM642 Data Sheet. (2005)
25. SMART Interconnect Group, U.: Flexsim 1.2 flit level simulator (2005)
26. Chen, X., Peh, L.S.: Leakage power modeling and optimization in interconnection networks. In: ISLPED '03: Proceedings of the 2003 international symposium on Low power electronics and design, ACM Press (2003) 90–95
27. Ho, R., Mai, K., Horowitz, M.: Efficient on-chip global interconnects. In: IEEE Symposium on VLSI Circuits. (2003) Stanford Univeristy.
28. Kim, J.S., Taylor, M.B., Miller, J., Wentzlaff, D.: Energy characterization of a tiled architecture processor with on-chip networks. In: ISLPED '03: Proceedings of the 2003 international symposium on Low power electronics and design, ACM Press (2003) 424–427
29. Stechele, W.: Algorithmic complexity, motion estimation and a vlsi architecture for mpeg-4 core profile video codecs. In: International Symposium on VLSI Technology, Systems and Applications. (2001)
30. Hendrickson, B., Leland, R.: The chaco user's guide, version 2.0, technical report sand94-2692 (1994) <http://www.ti.com/corp/docs/press/backgrounder/omap.shtml>.
31. Krashinsky, R., Batten, C., Hampton, M., Gerding, S., Pharris, B., Casper, J., Asanovic, K.: The vector-thread architecture. *SIGARCH Comput. Archit. News* **32**(2) (2004) 52
32. Marculescu, D.: Application adaptive energy efficient clustered architectures. In: ISLPED '04: Proceedings of the 2004 international symposium on Low power electronics and design, ACM Press (2004) 344–349
33. Zyuban, V.: Unified architecture level energy-efficiency metric. In: GLSVLSI '02: Proceedings of the 12th ACM Great Lakes symposium on VLSI, ACM Press (2002) 24–29
34. Hartstein, A., Puzak, T.R.: The optimum pipeline depth considering both power and performance. *ACM Trans. Archit. Code Optim.* **1**(4) (2004) 369–388
35. Fisher, J.A., Faraboschi, P., Desoli, G.: Custom-fit processors: letting applications define architectures. In: MICRO 29: Proceedings of the 29th annual ACM/IEEE international symposium on Microarchitecture, IEEE Computer Society (1996) 324–335
36. Heo, S., Asanovic, K.: Replacing global wires with an on-chip network: a power analysis. In: ISLPED '05: Proceedings of the 2005 international symposium on Low power electronics and design, New York, NY, USA, ACM Press (2005) 369–374
37. Taylor, M., Kim, J., Miller, J., Wentzlaff, D., Ghodrat, F., Greenwald, B., Ho, H., m Lee, Johnson, P., Lee, W., Ma, A., Saraf, A., Seneski, M., Shnidman, N., Frank, V., Amarasinghe, S., Agarwal, A.: The raw microprocessor: A computational fabric for software circuits and general purpose programs (2002)
38. Morad, T.Y., Weiser, U.C., Kolodny, A., Valero, M., Ayguade, E.: Performance, power efficiency and scalability fo asymmetric cluster chip multiprocessors. In: CCIT Technical Report 514, Technion (2005)