

Stochastic spatial routing for reconfigurable networks

André DeHon^{*}, Randy Huang, John Wawrzynek

University of California at Berkeley, Berkeley, CA 94720-1776, USA¹

Available online 28 February 2006

Abstract

FPGA placement and routing is time consuming, often serving as the major obstacle inhibiting a fast edit-compile-test loop in prototyping and development and the major obstacle preventing late-bound hardware and design mapping for reconfigurable systems. We introduce a stochastic search scheme which can achieve comparable route quality to traditional, software-based routers while being amenable to parallel, spatial implementation. We quantify the quality and performance of this route scheme using the Toronto Place-and-Route Challenge benchmarks. We sketch hardware implementations ranging from a minimal hardware-search assistance scheme which provides two orders of magnitude speedup, to FPGA-based schemes which provide greater speedup, to full hardware schemes which provide over three orders of magnitude routing acceleration. For coarse-grained devices with wide-word datapaths, the area overhead for integrating this hardware support into the network can be below 30%; for conventional FPGAs, a collection of hundreds of FPGAs can be configured to route one FPGA rapidly. With parallel path searches, the time required for the spatial solution scales sublinearly in network size for the typical, limited-bisection networks used for practical reconfigurable systems.

© 2006 Elsevier B.V. All rights reserved.

Keywords: FPGA; Detail routing; Reconfigurable computing; Spatial routing; Randomized algorithms

1. Introduction

Several researchers have attempted to reduce the time taken to route a user design by tuning software routing algorithms or attacking the problem with multiple processors. They were, at best, able to achieve route times on the order of seconds (see Section 4). For fast turn-around in the edit-compile-debug cycle and rapid prototyping, these software approaches still require billions of processor cycles and are not sufficient to make runtime routing viable in circumstances where it might benefit dynamic spatial computations, which require runtime mapping (Section 2.2).

At the core of almost all routing algorithms is a path search to find the least cost path through the network between a source and a sink. By adding a modest amount of support hardware to the routing network, we allow the

network itself to assist in the search for the least cost routing paths. Our solution uses this augmented network to find all available paths between a source–sink pair in time proportional to the distance of the route using a parallel, hardware search. Using this hardware-assisted technique, the time for the search task can be reduced 1000-fold over the software version. This suggests a large family of possible hardware–software routing solutions that range from minimal hardware to assist path search to full hardware to support route allocation and rip-up. The hardware scheme can be integrated into the network itself or be realized using an FPGA-based accelerator.

We start by motivating the demands for fast routing (Section 2). We review background and prior art (Sections 3 and 4). In Section 5, we outline the basic hardware-assisted routing scheme. This raises many issues which must be addressed to assemble complete routing schemes. We highlight these in Section 6, quantify the route quality in Section 7, and then go on to address hardware–software implementation options in Section 8. In Section 9, we analyze the performance of various implementation options and in Section 10, we address the hardware costs for the

^{*} Corresponding author.

E-mail addresses: andre@acm.org (A. DeHon), rhuang@cs.berkeley.edu (R. Huang), johnw@cs.berkeley.edu (J. Wawrzynek).

¹ <http://brass.cs.berkeley.edu/>.

various schemes. In Section 11, we briefly note that the techniques can be generalized for other kinds of configurable networks before concluding (Section 12).

Work in the 1980s (see Section 4) introduced the idea of hardware-assisted routing. However, that work was limited to supporting only the Lee’s algorithm path search kernel [24] for VLSI and printed-circuit design problems. We advance this pioneering work in several ways, making the solution more broadly applicable, addressing system- and application-level issues, and providing performance quantification. We:

- Adapt path search to the FPGA routing problem.
- Address routing beyond Lee’s path search, including state management, rip-up and retry, and fanout support (Section 6).
- Detail system-level solutions.
- Develop a variety of implementation alternatives, including augmenting custom hardware, hybrid hardware/software solutions, and FPGA-based routing engines (Section 8).
- Estimate the costs in custom and FPGA implementations (Section 10).
- Quantitatively assess the overall quality and performance impact of various implementations, including direct comparison with a modern software routing algorithm, Pathfinder [27] (Sections 7 and 9).

This article summarizes ideas introduced in [14,19], however the detailed quantitative comparison here is completely revamped and expanded from [19].

2. Motivation

2.1. Late bound designs

To take fullest advantage of reconfigurable, spatial computing platforms, we want to specialize the instantaneous computation to both the problem being solved (e.g., Instance-Specific SAT [37] and Covering [29]) and the available resources in the platform. Both of these may be *late bound* quantities. That is, for portability, the exact set of resources available on the platform may not be known until the program begins to run (e.g., SCORE Section 2.2), and the characteristics of a particular problem are not known until the program sees the problem. In cases such as these the time required to partition, place, and route the design onto the platform is part of the critical runtime of the application, not part of a pre-runtime compilation process. As such, any runtime taken to solve these tasks diminishes the potential acceleration offered by the spatial computing platform and narrows the domain of application where the reconfigurable platform or the specialized solution is superior to the more conventional alternatives.

Our goal in this respect is to understand how far we can compact the time required for routing of general-

purpose computational graphs (i.e., without exploiting task-specific route structure). Here, we focus entirely on routing. Wrighton and DeHon [36] demonstrate a spatial placement scheme that achieves similar performance improvements to the ones we demonstrate here for spatial routing.

2.2. SCORE and runtime mapping

We have been developing SCORE, a stream-based compute model which virtualizes reconfigurable computing resources (compute, storage, and communication) by dividing a computation up into fixed-size “pages” and time-multiplexing the virtual pages on available physical hardware [8]. SCORE’s goal is to serve as an abstract interface level, like an Application Binary Interface (ABI) or an Application Programmer Interface (API) in conventional programmable processors, that abstracts the detailed hardware implementation, including the number and kinds of resources, from the application and programmer. This allows SCORE designs to migrate automatically to newer, larger hardware platforms and exploit the additional resources. A key consequence of this abstraction is that the compiler does not know how big the SCORE platform will be. Consequently, the SCORE runtime must map the abstract SCORE graph onto the available physical hardware, time-multiplexing the large SCORE graph onto smaller hardware as necessary.

The SCORE runtime must perform routing no sooner than application load time and may need to perform routing as frequently as every reconfiguration in the time-multiplexed execution. Therefore, route time will reduce the raw performance potential of a SCORE application. If routing takes more time than a typical time slice, routing will have to be amortized across multiple time-slices to be viable.

The use of fixed-size compute pages connected by multi-bit busses reduces the size of the runtime placement and routing problem by a constant factor. This has the effect of simplifying the routing task, but does not make it trivial or address long-term scaling. As we will see in Section 10.1, the multi-bit busses and larger grained pages allow us to amortize the hardware overheads associated with hardware-assisted routing, making their cost quite small for the typical SCORE case. Our current impression is that a compute page should contain around 1000 bit-level operators (e.g., 4-LUTs or adder-bits) and busses will be 4–16 bits wide.

2.3. Route time scaling

The work required to solve our routing problems is scaling faster than sequential processor performance. Chip sizes, including FPGAs, scale with Moore’s Law; the number of nets in a design scales linearly with FPGA size, and the work required to find quality routes scales faster than linearly in the number of nets. At the same time, sequential

processor performance is scaling slower than Moore’s Law [1]. In separate experiments (see Chapter 5 of [20]), we found that if we limited the number of path searches (but not the work done per path search) to a constant multiple of the number of nets, the resulting, Pathfinder route quality decreased with increasing network size (i.e., the minimum channel width which Pathfinder could find increased compared to the optimum channel width), suggesting that the number of route trials required to find an equivalently good solution certainly increases faster than linearly in network size. Additionally, the work per route trial increases with larger networks, due both to the increased path lengths and the increased number of potential paths to search. Consequently, we are faced with the following choices:

- Route times continue to increase.
- Route quality decreases.
- We employ more scalable, parallel techniques, such as the one described here, to contain route time while maintaining route quality.

3. Background and definitions

3.1. HSRA

We build on the Butterfly Fat-Tree (BFT) [17] or equivalently, the linear switch population HSRA [33] (see Fig. 1). A key feature of this network is that the number of switches in each hierarchical switchbox is linear in the number of wires in the switchbox and the total number of switches in the network is linear in the number of endpoints [13].

This network has an important property which is not shared by Manhattan arrays: there is a unique set of switchboxes between any source and sink. Consequently, global routing is trivial (there is only the one solution), making detail routing our only concern. For our hardware-assisted router, this also means there is a unique “least common ancestor” or “crossover” switchbox between any source and sink (see Appendix B.1 for a formal definition). We use this localization to detect route success, or failure, locally in the crossover switchbox. Further, once we select a particular wire (switch) in a crossover switch-box, the path from the crossover to the source and sink, including the set of switches and wires in the path, is completely unique; this property simplifies path identification and allocation.

3.2. Pathfinder

Pathfinder is the dominant approach to FPGA routing currently in use in the academic community and heavily used in industry as well. For the limited switching available in typical, efficient FPGA and reconfigurable designs it is viable to fully represent the routing graph so that global and detail routing are combined into a single route

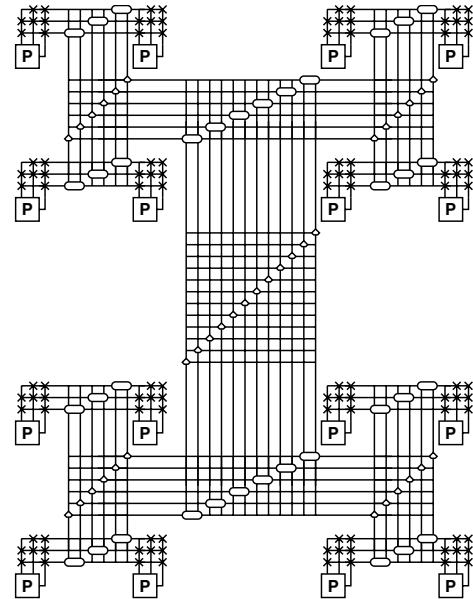


Fig. 1. HSRA network topology. P’s represent the network endpoints (e.g., LUTs or SCORE Compute Pages). The circles and ovals are switchpoints (see Fig. 2 for details); X’s show the connection-box switches. The network shown has three base channels.

problem; this resulting task is effectively an integer multicommodity flow problem which is known to be NP-Complete [15]. Pathfinder forms the basis of the previous, software-based attempts to accelerate routing (Section 4). While later versions of Pathfinder have been tuned for Island-Style FPGAs (e.g., [4]), the original Pathfinder formulation [27] is quite generic and can be applied to any routing network.

Starting from the base Pathfinder algorithm, we implemented our own versions for the HSRA [33,20]. We believe our implementation is very close in spirit to the original. The basic algorithm is as follows:

1. Create a fixed ordering of all nets in the design.
2. While there are unrouted nets and we have not exceeded the maximum number of route trials:
 - (a) For each net in the original fixed ordering
 - If net is unrouted (no path or shares paths)
 - Perform a *route trial* \equiv rip-up the congested net and reroute by performing a shortest-path search on the graph
 - (b) Increase history cost for each congested net

The basic routing step in Pathfinder is a shortest path search on the graph where each link is weighted by its congestion (over use); in this manner, the shortest paths will be the ones with least congestion, preferably paths with no congestions. This is the same basic approach taken in Linear Programming (LP) approximation algorithms for concurrent multicommodity flow (e.g., [25]), and the link costs can be viewed as dual variables in a primal–dual LP formulation.

To resolve congestion, Pathfinder adapts the cost of resources using a history term. If a resource remains congested, its history is increased and hence its cost is increased. Routes which can avoid the congested resource will be rerouted to do so. This tends to drive paths away from the congested region.

4. Prior work

Many attempts have been made to accelerate router performance, including optimizing the software, using parallel processors, and designing custom VLSI hardware.

Swartz, Betz, and Rose employ depth-first search and focused target selection to tune a Pathfinder-based router to decrease route time in low-stress routes (those where the router is allowed to use more channels than those required by the channel minimizing Pathfinder). They show that this combination leads to a router which requires roughly 1.1 ms per LUT/FF pair using a 300 MHz SPARCstation [31]. Normalizing for hardware technology, this means roughly 300,000 cycles per LUT/FF pair, or, assuming an average of four input nets per LUT-FF, about 75,000 cycles to route each two-point net.

Tessier [32] used domain negotiation and A* search to achieve similar reductions showing 200,000–500,000 cycles per 4-LUT. Lola, a fast, greedy, maze router, achieves roughly 500,000 cycles per net [16].

Our own fast, software router [20] achieves times similar to these examples, requiring roughly 140,000 cycles per net (see Table 3). Given the variations in machine architecture and increasing relative memory costs, cycles are a crude metric for comparison. Nonetheless, this establishes a consistent base range for our detailed hardware–software comparison.

Parallel processing is an obvious direction to further decrease route times. Chan and Schlag [9] attacked the problem of FPGA routing times with both coarse-grained parallelism and FPGA-accelerator assistance. They were able to show a 2–4× speedup in route time using 4–5 processors for congestion- and delay-driven [10] routing.

In the early 1980s, a number of researchers designed systolic array hardware to implement Lee’s Maze rout-

ing algorithm [24] including [7] (also known as *Pathfinder*) [18,21,30,35]. Our router approach is similar in spirit to these routers. We adapt the connectivity from a simple grid to FPGA switching networks, and we develop schemes for fanout, congestion negotiation, and victimization which go beyond the simple least-cost path search of these early VLSI routers. We further provide full application-level solutions and quantitative assessment.

5. Outline of basic solution

The key idea in our hardware-assisted router is to use the network structure itself to support the parallel path search and to keep track of the state of the network.

To find an available route in the HSRA network, we start at the source and the sink node and trace free (least cost) paths from the source and sink to the crossover switchbox. If the search from the source and the search from the sink meet on one (or more) wire(s) at the crossover switchbox, we have found a viable route path. We can then allocate the path (one of the paths) to this source–sink pair.

A pair of typical HSRA switchpoints is shown in Fig. 2. The switches allow us to make connections as appropriate, connecting the children for crossover connections, or connecting the appropriate child to the parent for up and down connections.

Now, consider adding a logical OR between the two children channels and placing the result on the associated parent channel (see Fig. 3). This additional OR allows an unused switchpoint to propagate a one (1), used to indicate a path search, up through unused parent links; the associated AND disables this OR-up functionality when the switch is already in use. With this addition, we can perform a *route trial* roughly as follows:

1. Set all endpoints (e.g., LUTs or SCORE compute pages) to drive zeros into all unused input and output connection to the network and all allocated source lines (leave allocated sink lines undriven as they will be driven by their associated sources).

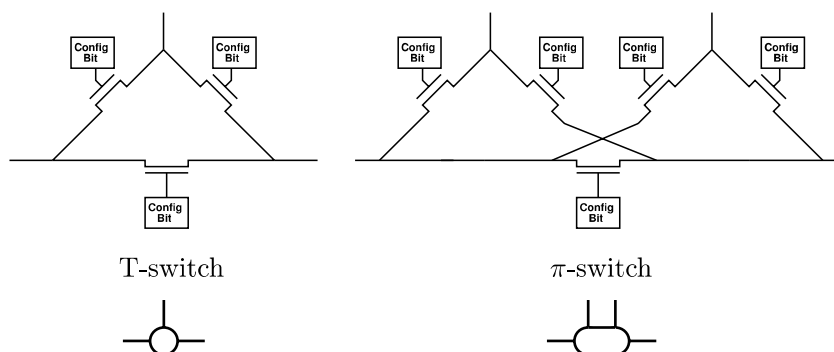


Fig. 2. HSRA switchpoints (pass transistor switch implementation).

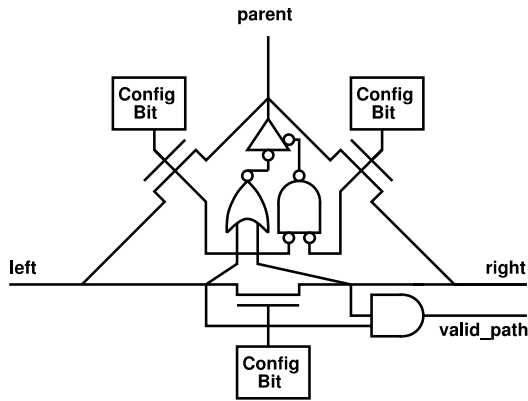


Fig. 3. HSRA T-switch with path-search or. The π -switchpoint can be augmented in a similar manner. This logical augmentation can be easily adapted for rebuffered and clocked switchpoints as well.

2. For the designated source–sink pair which we are currently trying to route, drive a one into each unused (available) network connection.
3. Wait for the driven ones to propagate through the network (through the OR’s as in Fig. 3) to the unique crossover switchbox.
4. At the crossover switchbox, scan for a switchpoint which receives a one on both of its sibling sides (see `valid_path` in Fig. 3); only this source–sink pair is driving ones, so a matched pair of ones indicates a complete path from both the source and the sink.
5. Allocate the unique path associated with one such matched pair; this means we go ahead and set the switches accordingly to connect this path. Note that this means this path will have zeros driven into it in the future and will not be considered in subsequent path searches.

Fig. 4 shows an example of this path search. We perform this search and allocation route trial successively for every

network connection in the design (e.g., Step 2(a) in the Pathfinder outline in Section 3.2).

The prospect for acceleration here is simple. In the traditional, software path search, each route trial takes several tens of thousands of cycles (e.g., see Table 3) to walk a network data structure and to explore all the possible paths between source and sink until a free (or inexpensive) path is found. In this hardware case, we use the network itself to explore all paths simultaneously. It does so quickly because all the switched paths are instantiated in hardware and directly connected by wires. It takes only the signal propagation delay across the wires and switches to trace back all possible paths. If the subsequent allocation can be performed cheaply in place, this turns the whole task from several tens of thousands of cycles into a just a few tens of cycles. Further, by either embedding this logic in the reconfigurable device or using collections of devices (e.g., many FPGAs), the spatial hardware acceleration scales with the hardware problem size (see Section 8.6).

6. Details

To obtain a complete scheme, we must fill in the details left open in the sketch above. In the next couple of section, we highlight these issues and provide detailed solutions. Here, we need to answer:

- How do we select among available paths?
- What do we do when no routes are found?
- How do we handle fanout? (Section 6.2)
- How do we perform allocation and victimization? (Section 8)

A key issue with respect to the traditional, software Pathfinder is history and costs. In the simple scheme above, we only have binary costs – either a path is free or it is not. Pathfinder, however, allows routes to share

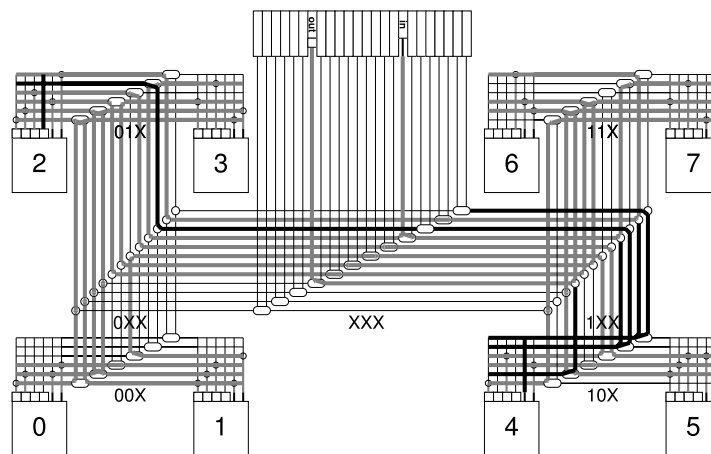


Fig. 4. Path search. Shown here is the result of a path search for a route from node 4 to node 2. The light, thick lines show pre-existing routes. The dark, thick lines show the paths driven to ones by the source and sink and propagated via the up or logic. At the crossover switchbox (labeled XXX), there is only a single switch which has a one arriving from both sides. We allocate the path that is joined by this switch. Note that there is a single, unique path from the source (node 4) to the sink (node 2) through this switch.

paths and uses congestion and historical congestion to bias the cost of paths. We show that suitable use of random path selection (Section 6.1) and coarse-grained route locking (Section 6.2) is a sufficient alternative to Pathfinder's history function.

6.1. Stochastic path selection

In place of history, we take advantage of randomness and our fast route exploration. When there are multiple free paths, we select randomly from among the available paths. That is, the parallel path search will explore all paths and find any free paths which exists. When there is no free path, we identify the set of least cost paths and select one path randomly from this path set; we rip-up the nets which interfere with the selected path and then allocate the path to the current route. Since, we may select a different path on each reroute of a net, the random path search performs its own sort of path negotiation. Further, since each full network route attempt fills in paths differently, we can make multiple attempts to route the full network from an empty network and select the best one. That is, since we select randomly among equal-cost paths, routes added early may or may not be placed on paths needed by later routes. In a history-based scheme, like Pathfinder, congestion negotiation adapts the cost of links to eventually migrate nets with choice away from critical links. Without history, we instead use random path allocation to perturb the network and explore different choices in the assignment of nets to routing resources. We call each full network route attempt a “try” and use the variable N_{try} to denote the number of such route attempts we perform to obtain a given route quality (i.e., this adds an outerloop to the Pathfinder algorithm in Section 3.2 where that loop is performed for each “try”).

6.1.1. CountNet path cost

When there are no free paths for a route, our path search selects a path to allocate anyway. The nets which must be ripped-up to allow the new path to be allocated are victimized (Sections 8.4 and 8.5). There are several ways we could select the path to allocate in this case. After experimenting with several alternatives including completely random path selection and minimum congested segment (CountCongestion) rip-up selection [19], we have settled on a scheme that selects a path randomly from the set of paths which will disrupt the least number of existing nets.

Ideally, we count the number of nets that would be victimized if each path were selected. In Fig. 5, we show one

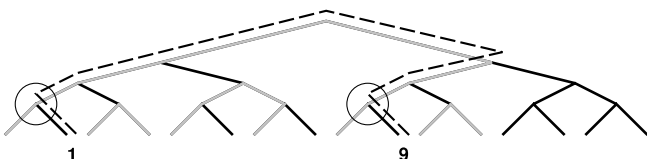


Fig. 5. Path search with CountNet heuristic.

wire channel of a size-sixteen ($p = 0$) tree. Lightly shaded wire segments are occupied by existing nets; unoccupied wire segments are shown in black. Suppose we were to perform a path search from node 1 to node 9 (shown with a dotted line). This net will disrupt six wire segments, but will only victimize a single, pre-existing net. This “CountNet” scheme directly reflects the number of existing nets affected by this path and, consequently, the amount of re-routing work that has to be done when a path is chosen to be ripped-up.

6.1.2. Represent route costs

We can add costs to our path search by delaying the path search signal. That is, we pass free (e.g., uncongested) path search signals through to the parent node with no additional delay. When a path search signal encounters switches which are in use, it is delayed for a cycle and then passed up (see multiplexer in Fig. 6). In this way, the first set of path search signals to reach the crossover switchbox will represent the least cost paths; strictly speaking, since the search from the source and sink occur in parallel, the first set of crossover switches which see a path search from both sides are the one's where the maximum cost of the source path and the sink path is minimized. See Appendix B for precise statement.

6.1.3. Implementation approximation

We use this delayed route signal scheme to cheaply represent and count the cost of each path. However, the ideal CountNet cost heuristic could be expensive to implement in hardware as it requires us to store a net ID and perform a comparison at every switch. To approximate the CountNet heuristic, we observe that a search signal and a routed net can “interact” at only two switches (entering and existing) as shown in Fig. 5 (black circles). At those two switches, the configuration of the switch will be different from the search direction. If a switch is occupied and has a configuration different from the search direction, we have encountered a new net and will delay the search signal by one cycle. As we can see in Fig. 6, the CountNet approximation implementation requires roughly 10 gates. Our measurement shows that this scheme will, on average,

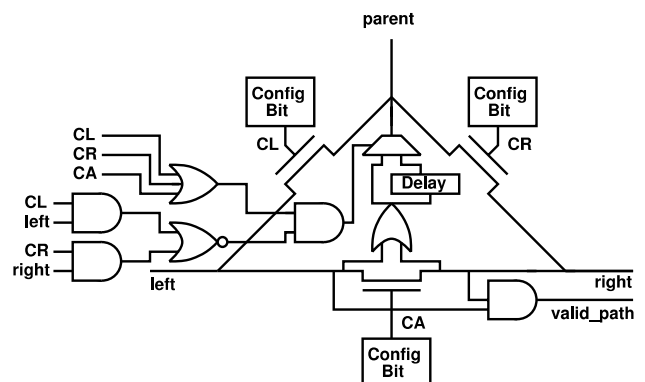


Fig. 6. HSRA T-switch with CountNet approximation delay.

choose the same net as counting the exact number of nets a path has to victimize 75% of the time [20]. Appendix B provides a more formal definition of these cost functions.

6.1.4. CountNet heuristic route quality

To understand the quality of this approximation, we generated a set of difficult synthetic benchmarks to ensure our solution performs reasonably with larger and harder designs; these benchmarks make sure to maximally fill many switchboxes according to the switch-box population, assuring that we can differentiate the effects of route quality and limited population switch-boxes (see Chapter 5 of [20]). We scale the synthetic network size from 8 to 4096 and, for each size, we generate 100 netlists which stay unchanged throughout our experiments.

In Fig. 7, we show the quality of the ideal CountNet scheme along with the approximation; we allow both schemes to select the best result from 20 route starts ($N_{\text{try}} = 20$). Here, quality is measured in the number of base channels (C) required to route the design; more channels indicate lower quality. We also include results from the history-based, software Pathfinder as a base-line comparison. For the history-based Pathfinder, we route each of the 100 netlists once and then average the results for each array size. We use a route trial multiplier (RT) of 100 for the history-based Pathfinder; this means we allow the router to attempt a number of individual, two-point net route trials equal to $100 \times$ the total number two-point nets in the design. We see that the CountNet approximation heuristic requires no more channels than the history-based Pathfinder up to the 1024 node networks; above 1024 both the ideal CountNet scheme and history-based Pathfinder achieve comparable quality which is only slightly better than the approximation.

For the smaller array sizes, the randomization allows count net schemes to explore route sets which Pathfinder does not; while Pathfinder explores all individual routes, it does not explore all aggregate combinations. That is, because of the randomization, if we performed enough trials, the stochastic router would eventually explore all possible route sets. If the problem is routable, there is an assignment of nets to paths that routes the design. Each such assignment can be seen as a choice for each net among

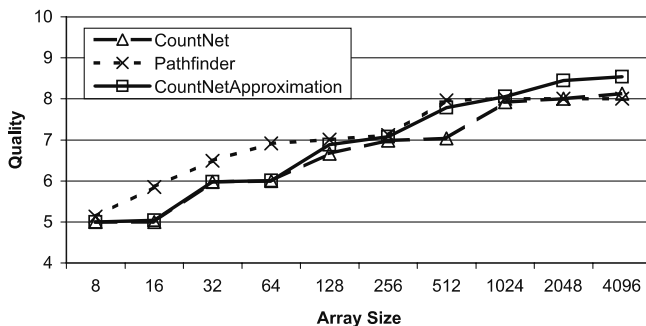


Fig. 7. Count cost vs. Pathfinder comparison. Both CountNet schemes use 20 route starts.

the free paths. So, if we explored all choices for free paths, we would explore all routes. The question is how many such trials would we need to make in order to explore all possible paths? This is like the “infinite number of monkey’s typing on a keyboard” scenario where they will eventually produce the works of Shakespeare [11]. Certainly, it would not be viable to perform enough route trials to guarantee all route sets were explored for networks of any reasonable size; nonetheless, this thought experiment illustrates that the approach can eventually find solutions if they exist. Traditional Pathfinder, on the other hand, is entirely deterministic and is known not to converge in some situations [10].

6.2. Hypergraph support

To route nets with fanout (i.e., hypergraphs), we sequentially route each two-point net, trying to re-use as many links as possible from existing paths allocated to this net. This allows us to avoid the over-constraint and complexity of dealing with multiple sinks simultaneously.

This is similar to the traditional, software-based Pathfinder which also routes two-point nets as the primitive route step. In the software case, the algorithm keeps track of which net(s) are actually using a resource so that it can allow two-point nets driven by the same source to share the resource and count as a single net.

6.2.1. Fanout routing

We add a state bit at every switch to keep track of the two-point links which belong to the same net. This bit is set when we allocate the switch during the current net search. This bit is cleared when we begin to route a new net.

We also add a *global route tree* which is used for both fanout routing and allocation (Section 8.3). The global route tree is a binary tree that looks just like the OR-up logic in the T-switches, except that it has no configuration bits; since this is a binary tree, the global-route tree has only a single T-switch in each switchbox and is, therefore, only a small additional cost on top of the large number of switchpoints and wires already in each switchbox. At the leaves, this tree, which is separate from the normal routing paths, is also driven by the source and sink participating in the path search.

We implement the following simple scheme:

1. Order the destinations associated with a single source by decreasing path length. The path length is twice the height of the crossover switchbox in the HSRA.
2. For each destination:
 - From the sink, we send a search signal on all unused inputs and drive the global route signal.
 - From the source, we do nothing and do not drive the global route signal.
 - At a switch the global route signal tells us which direction is the sink side; the sink side will have the global route signal driven. The state bit helps us determine if the switch has been allocated during the current net

search and therefore can be a point of fanout for the current fanout search. If the state bit is set and the search signal from the sink side arrives without encountering congestion, we have found an available path.

- Otherwise, drive ones into all available source paths and allocate a new path, like a standard path search.

Since this scheme greedily selects paths one two-point link at a time, it will not necessarily find the minimum cost path set for a net. Nonetheless, it definitely uses fewer resources than treating each source–sink connection as a separate net.

6.2.2. Fanout lock approximation

In the current CountNet approximation heuristic, a net with 1000 fanouts will cost the same as a net with no fanout; if a net with large fanouts is victimized, a large number of two-point nets will be ripped out when the large fanout net is re-routed, resulting in slower convergence and worse route quality. To deal with this problem, we could count the number of fanouts that would be affected and choose the path with the least fanouts. However, implementing count fanout exactly in hardware appears prohibitively expensive.

We can approximate the count fanout heuristic in a binary fashion with a fanout lock. The idea is that we want to lock down nets with large fanouts after they have been routed and prevent them from being ripped-out. Effectively, this scheme says the cost of a victimizing the high fanout net is infinite so it should not be a victim candidate. Since we order nets by decreasing fanout, high fanout nets will be routed first before they have a chance to interfere with each other. To implement fanout lock in hardware, we:

- Add a lock bit for every switch.
- Assert the lock bit after allocation for high fanout net.

If a switch has an asserted lock-bit, it will not propagate cost signal upward. This guarantees that the crossover switch box will not select a path with high fanout nets. On sample graphs, we obtained our highest quality results when we set the fanout lock threshold between five and 10. If we set the fanout lock threshold too low, the router degenerates to a greedy router; the lower threshold of five is slightly above the average fanout in these 4-LUT graphs, guaranteeing that most nets are routed normally. If we set the threshold too high, then we may spend too many cycles victimizing and rerouting large nets. For the results which follow, we used a fanout lock threshold to 10.

7. Quality comparison

For comparison, we use the standard FPGA place-and-route benchmark suite from Toronto [5]. The benchmark suite is placed on a tree network using tools we developed for the HSRA [12]. The growth rate of the network is

governed by the growth exponent (p) in Rent's Rule [23] and is selected to be 0.6 based on our previous work [12]. Rent's Rule is an empirical relationship which characterizes the locality structure of a graph; that is, if we cluster groups of N graph nodes to minimize the number of nets which must enter or leave that collection of nodes, then the number of nets crossing in and out of the cluster (IO) will be:

$$IO = cN^p \quad (1)$$

The same placement is used for both routing algorithms and thus the comparison is fair regardless of the placement tool.

For the fanout lock algorithm, we route each netlist 500 times and use the statistics to calculate the expected number of tracks if we were to pick the best of multiple starts. For the Pathfinder algorithm, we set the route trial multiplier to 50 (RT), which represents a high routing effort.

We allow the number of base channels (C) (see Fig. 1) used for the route to float and measure quality in terms of the number of base channels an algorithm requires to route the design. In Table 1, we list the quality results of routing each benchmark using both algorithms. We see that as we increase the number of route starts (N_{try}), the chance of finding a better quality route increases as well. The results show that a random algorithm with no history, and hence simple enough to be implemented in hardware, can approach the history-based Pathfinder algorithm in terms of quality (within 3%).

8. Basic hardware and software support

In Sections 5 and 6, we introduced the hardware logic needed to perform path search. In this section, we describe a variety of hardware–software tradeoffs to support the remainder of this routing scheme and relate this to the time required to perform a route.

8.1. Routing time

The time required to route a netlist with hardware assist is:

$$\begin{aligned} T_{\text{netlist}} = & N_{\text{try}}(N_{\text{RT}} \cdot (T_{\text{ctrl}} + T_{\text{path}} + T_{\text{check}} + T_{\text{alloc}}) \\ & + N_{\text{RO}} \cdot T_{\text{victim}} + N_{\text{FO}} \cdot (T_{\text{ctrl}} + T_{\text{path}} + T_{\text{check}}) \\ & + N_{\text{FOA}} \cdot T_{\text{alloc}}) \end{aligned} \quad (2)$$

As previously noted, N_{try} is the number of route starts, N_{RT} is the total number of route trials (Section 5), N_{RO} is the total number of ripouts, N_{FO} is the number of augmenting fanout routes, and N_{FOA} is number of allocations following successful fanout searches. T_{ctrl} is the number of cycles to send a control signal. Control signals travel from the root to the leaf nodes in $O(\log(N))$ cycles. T_{path} is the number of cycles it takes to propagate a signal from the leaf nodes to the crossover switch box (i.e., steps 1–3 in Section 5).

Table 1
Route quality comparison

Design	LUTs	Size	N_{try} for spatial route			Pathfinder RT = 50
			5	10	20	
alu4	1522	4096	10.00	10.00	10.00	10
apex2	1878	4096	10.59	10.35	10.12	11
apex4	1262	4096	11.00	11.00	11.00	11
bigkey	1707	2048	8.27	8.07	8.01	9
clma	8383	32,768	11.00	11.00	11.00	11
des	1591	2048	10.00	10.00	10.00	9
diffeq	1497	2048	9.96	9.92	9.84	8
dsip	1370	2048	8.57	8.32	8.10	9
elliptic	8192	8192	10.04	10.00	10.00	10
ex1010	4598	8192	13.22	12.97	12.80	10
ex5p	1064	4096	11.00	11.00	11.00	10
frisc	3556	8192	10.82	10.67	10.45	10
misex3	1397	4096	10.21	10.04	10.00	11
pdc	4575	16,384	12.01	12.00	12.00	12
s298	1931	4096	9.49	9.24	9.06	9
s38417	6406	8192	10.00	10.00	10.00	9
s38584.1	6446	8192	9.00	9.00	9.00	9
seq	1750	4096	10.05	10.00	10.00	11
spla	3690	16,384	12.57	12.33	12.11	12
tseng	1047	2048	10.00	10.00	10.00	8
Total			207.81	205.92	204.50	199

Spatial search uses the fanout lock approximation. Size indicates the size of the physical tree on which the benchmark is placed; these are all powers of two and may be depopulated to match the $p = 0.6$ growth rate [12].

T_{path} is net dependent and bounded by $O(\log(N))$. For modeling, we assume:

$$T_{\text{ctrl}} = \log_2(N) \quad (3)$$

$$T_{\text{path}} = \log_2(N) \quad (4)$$

In practice, the geometric path length mean is about 77% of this length; we use the Eq. (4) length so our estimates are conservative. T_{check} is the time it takes to generate a random number and check for available routes at the crossover switch box (Section 8.2). T_{alloc} is the number of cycles it takes to allocate a route (Sections 8.3 and 8.5). T_{victim} is the number of cycles it takes to victimize a channel at the crossover switch (Sections 8.4 and 8.5).

8.2. Random path selection

We can perform random path selection economically in the switchbox by using a pseudo-random number generator (PRNG) and a cyclic segmented parallel prefix (CSPP) circuit [22]. The PRNG indicates which crossover switch is preferred for allocation. We mask out those which are not selected, and use the CSPP circuit to identify the first circuit candidate switchpoint identified by the path search. The CSPP circuit allows us to identify the path in $O(\log(W))$ time, where W is the number of switchpoints in the switchbox. W grows as $O(N^p)$ ($1.0 > p > 0.5$; p is the exponent in Rent's Rule (Eq. (1)) which can be used to characterize the growth rate of bisection bandwidth in the HSRA [33]), so the depth of the CSPP circuit grows only as $O(\log(N))$. Our experi-

mental results using this PRNG-CSPP random number generation scheme are statistically indistinguishable from results generated using a pure random number generation. We assume:

$$T_{\text{check}} = \left\lceil \frac{\log_2(N)}{4} \right\rceil + 1 \quad (5)$$

8.3. Hardware allocation

We can build a route allocation mechanism into the network with an AND gate for each switch and an “allocate” pull up at the crossover; hardware allocation uses the global route tree which we have already introduced in Section 6.2.1 (see Fig. 8). Once the path search has found a possible path, we stop driving the normal network paths and drive an “allocate” request, a one, back down the selected path to perform the allocation. Each switch which receives this one performs the actual allocation on the appropriate parent-child link, propagating the allocation, in turn, down to that child; the global-route tree shows the switchpoint which child connection to allocate.

Note that, if we simply tried to allocate from the top without the global-route tree, the switchpoint would not know which child connection to make; the global-route tree provides this information. A similar problem occurs if we try to allocate from the bottom; without additional information, it is not clear which of the two up connections in a π -switchpoint the route should allocate.

During the allocation phase, we send the three control signals in sequence. These control signals can be pipelined,

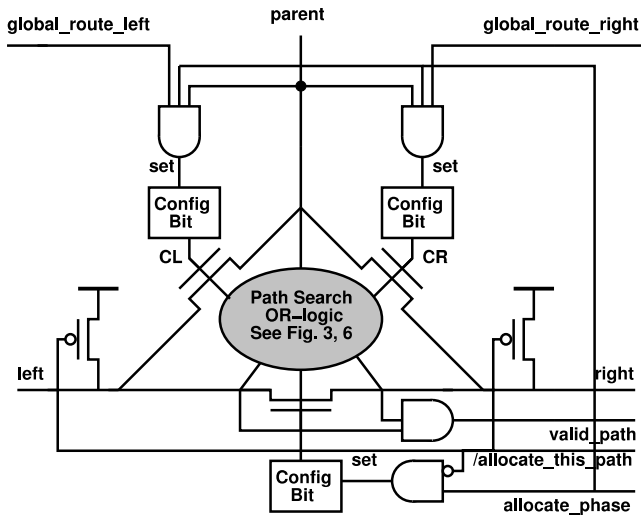


Fig. 8. HSRA T-switch with allocation logic. Allocate logic added to pass-transistor T-switchpoint: at the termination of path search, we do 3 things in sequence: (1) deassert the source and sink drive, but leave the global route tree driven, (2) assert allocate_phase to enable allocation, and (3) assert allocate_this_path to drive a one into the selected path at the crossover switchbox. Primed in this manner, the single path selected at the crossover is allocated along successively stages all the way down to the source and sink.

so T_{alloc} equals one trip through the network plus a small constant number of cycles to drain the pipeline. With this hardware allocation:

$$T_{alloc} = \log_2(N) + 2 \quad (6)$$

8.4. Hardware victimization

When there are no free routes found, we need to deallocate (victimize) existing routes in order to make a new route. Complete logic to support this in hardware is shown in Fig. 9. The logic needs to identify the intersecting paths and propagate the fact that the path is a victim to all switchpoints along the path before actually clearing the switchpoints. All together, this takes three crossover-to-leaf trips in the network to clear routes plus a 4th trip to perform the new allocation.

We also need to know which routes were victimized. At the end of the victim propagation, the sink will know, by the position of the input, which source it lost. If the sink knows which source is associated with this input, that is enough information for it to inform the route controller which source–sink pair(s) has been ripped up and needs to be re-routed. It is possible that many paths are victimized during a single deallocation. A binary collection tree would allow us to identify all victim paths in at most a number of cycles: $O(\log(N) + \text{number of victims})$.

During the steps of the victimization process, we have to wait for one step to finish before we can start another. Therefore, T_{victim} includes four trips through the network

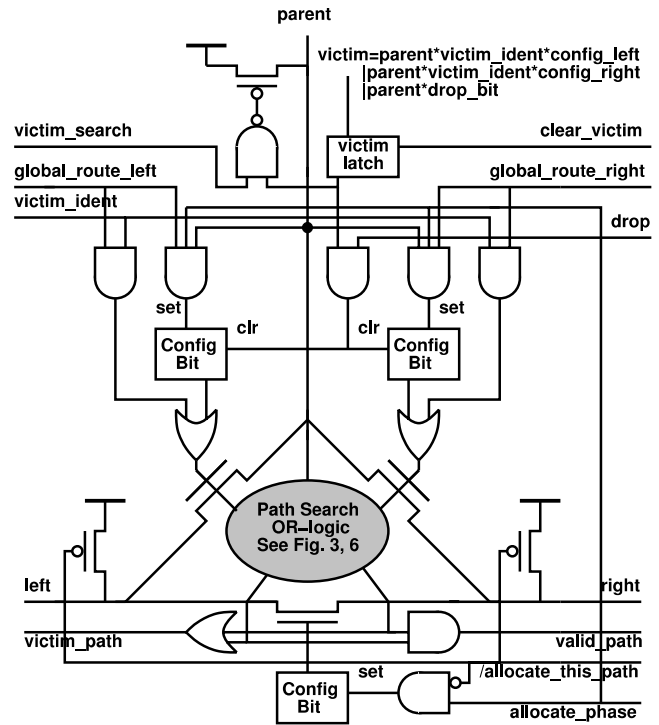


Fig. 9. HSRA T-switch with victimization logic. Victimization logic added to pass-transistor T-switchpoint: When we must steal a path in use, we do the following: (1) assert victim_ident and drive a one into the selected path at the crossover switchbox; this sets the victim latch everywhere the new path intersects an old path. (2) propagate the victim information up to the top of the victimized routes by asserting victim_search; (3) drive a one onto the victim paths from the crossover and assert drop_bit to mark all affected paths; (4) assert drop to clear the old paths; (5) poll the leaves to discover victims; (6) perform a normal allocation of the now cleared path.

plus the number of victims (V) plus a constant for signal sequencing.

$$T_{victim} = 4 \cdot T_{path} + V + 3 \quad (7)$$

8.5. Software allocation and victimization

A more modest solution, both in hardware cost and performance potential, is to perform only the path search in hardware (Fig. 3 or 6) and perform all record keeping in software. With no allocation or deallocation logic, we will need configuration bit addressability to set and clear configuration bits as routes are allocated and victimized.

The software bookkeeping needs a table to track switch usage. This table is indexed by switchpoint identification and contains:

1. Switchpoint identification for the left child switch of this switchpoint.
2. Switchpoint identification for the right child switch of this switchpoint.
3. The net that is using this switchpoint (if used).
4. Switchpoint identification for the crossover switch for the route through this switchpoint when in use.

This tables is $O(N)$ in size and is potentially accessed quite irregularly. It is unlikely this table will be cacheable, making each table reference moderately expensive.

Path search is initiated as before using the OR-up logic. When we find an available route, the controlling processor queries the crossover switchbox to discover which path was selected. The processor then walks the switchpoint table, starting at the crossover switchpoint, and stores the net identification and crossover switchpoint in each switchpoint along the path. As each switchpoint is visited, the processor issues a command to the network to allocate the appropriate switch bit.

When it is necessary to victimize paths, we perform the route victimization in software. Once we pick a target path, we:

1. Walk the switchpoint table starting at the victim crossover switchpoint.
2. For each switch visited we need to:
 - Find out if a net already occupies that segment,
 - If a net occupies the segment,
 - (a) Add that net to the list of unrouted nets
 - (b) Lookup the switchpoint which is the root of the existing net connection
 - (c) For each switchpoint belonging to this net
 - Issue a command to the network to clear the switch bit
 - Clear the net from the wire segment’s entry in the segment table

After deallocating the victim paths, we proceed with allocation.

In the allocate phase the processor reads one switchpoint entry from the switchpoint table for each switchpoint in the path. A path that crosses over at height, L_{path} , will have:

$$N_{\text{pathsw}} = 2 \cdot L_{\text{path}} - 1. \quad (8)$$

Like T_{path} , L_{path} is, at most, $\log_2(N)$, and, on average 77% of $\log_2(N)$. While each entry in the table will likely fit nicely in a cache line, access to the switchpoint table will almost certainly generate a cache miss and have to be satisfied from main memory. Therefore, we expect to pay a single main-memory reference time (T_m) for each of the N_{pathsw} lookups, followed by 1–3 cache reference times (T_c) when we read more than one word from a switchpoint entry. There are N_{pathsw} writes to set bits in the array (T_a) and N_{pathsw} writes to the wire segment table. These writes are to items which are likely to be in the same cache line and are potentially pipelineable using a write buffer (T_{wb}). Therefore, allocation is likely to take:

$$T_{\text{alloc}} = N_{\text{pathsw}} \cdot (T_m + T_c + T_{\text{wb}} + T_a) \quad (9)$$

For our calculations, we use:

$$L_{\text{path}} = \log_2(N) \quad (10)$$

As in the hardware victimization case, an allocated route will conflict with a number of victims, V . Noting that

deallocating a victim requires essentially the same operations as allocating a path, the deallocation takes roughly:

$$T_{\text{victim}} = N_{\text{pathsw}} \cdot (T_m + T_c) + V \cdot T_{\text{alloc}} \quad (11)$$

Let us assume $T_c = T_{\text{wb}} = T_a = 1$ cycle. If we assume $T_m = 50$ cycles, a 16K-node network takes: $T_{\text{alloc}} = (2 \cdot 14 - 1) (53) = 1431$ cycles. From our experiments, typical values of V are between 3 and 4 for large networks, so we will assume $V = 4$ here, giving: $T_{\text{victim}} = 27 \cdot 51 + 4 \cdot 1431 \approx 7100$ cycles. Maintaining these data structures in memory clearly becomes the dominant time cost if we go with this hybrid hardware–software scheme. For a modern, large-scale, FPGA, we would likely use on-chip memory such as the embedded DRAM block designed for SCORE/ HSRA [28]. Random access in this memory takes 14 logic cycles ($T_m = 14$) making $T_{\text{alloc}} \approx 460$ cycles and $T_{\text{victim}} \approx 2250$ cycles. These numbers should be compared with $T_{\text{alloc}} = 16$ and $T_{\text{victim}} = 63$ for the hardware case in the previous two sections.

8.6. Parallelism

Only the searches to the same top, crossover switchboxes need to be sequentialized. A path search in the left half of a network can proceed completely in parallel with a path in the right half of a network. In general, if the least common ancestors of the nets’ source–sink pairs are in different subtrees, the searches can proceed in parallel. This means, we can search sequentially for paths which crossover in the topmost switchbox then search in parallel for the paths which crossover in its immediate left and right switchbox. This parallel decomposition continues in turn. As a result, the ultimate sequentialization in this scheme is the sum of the maximum number of paths crossing over at each switchbox level rather than the total number of nets. For a typical network or design with $1.0 > p > 0.5$, this means the sequentialization goes as N^p rather than N . Ultimately, of course, we will saturate the processor’s time to give attention to starting and completing routes. For large designs, we might consider allocating a control processor to subtrees at some level(s). For large designs, there are many other good reasons to consider this in the SCORE case, at least. Consequently, it is possible for the entire route time to scale only as $O(N^p)$.

8.7. Example route times

To see how these numbers come together, we consider routing the `alu4` benchmark. This design has 1522 LUTs and we place it on a 4096 LUT HSRA. The history-based, software Pathfinder requires over 13 billion cycles to route the design at $C = 10$. Using our stochastic scheme, we measure $N_{\text{RT}} = 7654$, $N_{\text{RO}} = 2744$, $N_{\text{FO}} = 10574$, $N_{\text{FOA}} = 7127$, and $V = 3.4$. To guarantee over a 99% chance of achieving $C = 10$, we use $N_{\text{try}} = 3$.

8.7.1. Hardware assist

With hardware for the search logic, $T_{\text{ctrl}} = T_{\text{path}} = \log_2(4096) = 12$. Using the PRNG and CSPP circuit, $T_{\text{check}} = \frac{\log_2(4096)}{4} + 1 = 4$ cycles. With hardware support for allocation and victimization, $T_{\text{alloc}} = \log_2(4096) + 2 = 14$ cycle and $T_{\text{victim}} = 4 \cdot 12 + 3.4 + 3 = 54.4$ cycles. Putting it all together:

$$\begin{aligned} T_{\text{netlist}} &= 3 \cdot (7654 \cdot (12 + 12 + 4 + 14) + 2744 \cdot 54.4 \\ &\quad + 10574 \cdot (12 + 12 + 4) + 7127 \cdot 14) \\ &\approx 2,600,000 \end{aligned}$$

Compared to the Pathfinder result of 13B cycles to route this netlist, this is a speed up of over 5000.

8.7.2. Optimization: overlap netlist routing

A small improvement to the base hardware case can be made by overlapping independent operations. In the final phase of routing, we perform allocation; at the same time, we can start issuing commands to route the next net. By the time the routing commands arrive at the leaf nodes, the previous net's allocation step has completed. This optimization eliminates the T_{ctrl} term from the equation.

$$\begin{aligned} T_{\text{netlist}} &= 3 \cdot (7654 \cdot (12 + 4 + 14) + 2744 \cdot 54.4 + 10574 \\ &\quad \cdot (12 + 4) + 7127 \cdot 14) \\ &\approx 1,950,000 \end{aligned}$$

We achieve a speed of over 6700.

8.7.3. Optimization: parallelize netlist routing

As described in Section 8.6, the achievable speedup depends on how much we have to sequentialize netlist routing. The ultimate sequentialization is the sum of the maximum number of paths crossing over at each switch box level. For the `alu4` design, there are 1536 nets; the maximum number of nets which cross over at each switchbox level is 1 2 4 5 4 8 16 31 44 89 252 123 22; these sum to 601. So, this gives us the potential, additional, parallelism for routing non-overlapping nets of $\frac{1536}{601} = 2.56$. Exploiting this parallelism, the route requires roughly $\frac{1,950,000}{2.56} \approx 760,000$ cycles, for a speedup of over 17,000.

8.7.4. Software allocation and victimization case

As described in Section 8.5, assuming $T_m = 50$ and using Eqs. (9) and (11), $T_{\text{alloc}} = (2 \cdot 12 - 1) \cdot (53) = 1219$ and $T_{\text{victim}} = (2 \cdot 12 - 1) \cdot (51) + 3.4 \cdot 1219 = 5317.6$. We route the design in:

$$\begin{aligned} T_{\text{netlist}} &= 3 \cdot (7654 \cdot (12 + 4 + 1219) + 2744 \cdot 5317.6 + 10574 \\ &\quad \cdot (12 + 4) + 7127 \cdot 1219) \\ &\approx 98,700,000 \end{aligned}$$

This gives us a speedup of over 130 compared to the full software Pathfinder.

9. Performance analysis

In this section, we compare the performance of the various hardware and software schemes introduced. We do this by measuring the cycles for both the history-based software, Pathfinder and a software version of our stochastic path search algorithm. From the software version of our algorithm, we extract the key parameters (N_{try} , N_{RT} , N_{RO} , N_{FO} , N_{FOA} , V) necessary to calculate the performance of the hardware-assisted schemes.

To measure running time of the software router, we use the 64b TSC (Time Stamp Counter) timer on the processor to measure time in cycles. Our software router is written in C and compiled with the GNU C Compiler (version 2.95.2) using `-O3` option. Benchmarks are run on a 1.4 GHz Pentium 3-based system running Linux 2.4.18 with a 133 MHz system bus and variable amount of main memory (from 256 MB to 2 GB). In general, since we are performing graph traversals on a large data structure, most memory accesses will be cache misses. However, we make sure our entire data structure will fit in the main memory, so that we are not measuring performance derated by virtual memory thrashing.

For the history-based, software Pathfinder algorithm, we run each benchmark three times and report the minimum time. In Table 3, we list the minimum number of base channels needed to route each benchmark (C_{min}), the routing time for C_{min} , $C_{\text{min}} + 1$, and $C_{\text{min}} + 2$, and average number of CPU cycles needed to route a 2-point net (for $C_{\text{min}} + 2$). An “N/A” entry indicates that the router fails to route at that given number of tracks. Our implementation of the software, history-based Pathfinder algorithm averages 140K cycles per 2-point net, which is comparable with previous work (Section 4).

Table 2 summarizes the path search parameters for the fanout lock algorithm when targeting the minimum C achievable by both algorithms. With these we calculate the speedup with hardware assistance using Eq. (2). In Table 3 we list the routing times for C_{min} , $C_{\text{min}} + 1$, and $C_{\text{min}} + 2$ which are achievable over 99% of the time. We see that our stochastic, hardware fanout lock algorithm is able to achieve equal or better quality results on fourteen of the benchmarks. Summing the minimum channel requirements across all 20 benchmarks, the history-based, software Pathfinder requires 199 channels, whereas the fanout lock algorithm requires 202.

Table 4 uses the data from Table 2 and summarizes the speedup under the variations introduced in Section 8.

10. Cost analysis

10.1. Hardware augmentation

The minimum software allocation and victimization scheme requires only 3–4 gates per switchpoint (Fig. 3). This should be compared to the three configuration bits and pass transistors required for a minimum switchpoint implementation. The additional 3–4 gates are likely to

Table 2
Statistics from fanout lock router on Toronto benchmark set at minimum base channels achievable by both routers

Design	C	#Nets	N_{try}	N_p	N_{RT}	N_{RO}	N_{FO}	N_{FOA}	V
alu4	10	1536	3	2.56	7654	2744	10,574	7127	3.4
apex2	11	1916	4	3.1	2546	167	5130	4664	3.1
apex4	11	1269	2	1.77	1732	110	3452	3097	3.2
bigkey	9	1935	2	3.38	2398	143	4650	4329	2.9
clma	11	8443	2	2.93	45,608	14,695	71,088	48,173	4.3
des	10	1847	7	2.47	2956	360	5240	4479	3.7
diffeq	9	1560	539	2.83	2091	177	4008	3652	3.2
dsip	9	1598	2	2.94	2015	138	4259	3979	2.5
elliptic	10	3734	8	2.83	5419	838	9329	8476	2.0
ex1010	12	4608	534	4.01	5375	55	11,587	10,875	3.1
ex5p	11	1072	2	1.64	2918	680	4967	3786	4.0
frisc	10	3575	117	2.83	32,239	11,942	40,776	22,984	3.6
misex3	11	1411	2	2.22	1846	102	3809	3474	3.5
pdc	12	4591	6	2.17	6352	425	13,415	12,055	2.9
s298	9	1934	33	3.02	2702	710	5056	5008	1.1
s38417	10	6434	2	7.01	8506	596	15,651	14,167	3.8
s38584.1	9	6483	6	5.93	8117	607	15,383	14,354	3.1
seq	11	1791	2	2.06	2250	95	4574	4208	2.8
spla	12	3706	42	2.58	5121	371	11,085	10,029	3.6
tseng	10	1098	2	2.47	1926	310	3384	2863	3.8

N_p is the amount of parallelism available when we route non-overlapping nets in parallel. N_{try} is chosen to guarantee 99% likelihood of achieving C shown; expected route trials can be a factor of 6 smaller. N_{RT} is the number of path searches performed, N_{RO} is the number of victimizations, N_{FO} is the number of fanout searches, N_{FOA} is number of allocations following successful fanout searches, and V is the average number of victims per rip-up.

Table 3
Pathfinder and fanout lock route time comparison

Design	C_{min}	Software Pathfinder				Hardware route			Speedup
		Cycles (Millions)			cyc/2pt net	Cycles (Thousands)			
		C_{min}	$C_{\text{min}} + 1$	$C_{\text{min}} + 2$	@ $C_{\text{min}} + 2$	C_{min}	$C_{\text{min}} + 1$	$C_{\text{min}} + 2$	
alu4	10	13,173	799	769	142K	2600	243	232	5000
apex2	10	N/A	676	529	79K	69,087	1300	584	510
apex4	11	12,256	417	389	87K	437	195	150	28,000
bigkey	8	N/A	11,146	931	142K	10,135	556	255	20,000
clma	11	42,061	11,614	11,481	169K	26,670	1656	1619	1500
des	9	95,352	370	356	58K	N/A	656	246	560
diffeq	8	367,254	49,284	486	90K	N/A	130,502	418	370
dsip	8	N/A	4717	1662	285K	20,711	496	227	9,500
elliptic	10	924,960	3538	2224	182K	5,707	585	574	160,000
ex1010	10	1,529,242	318,979	62,709	292K	N/A	N/A	412,658	150
ex5p	10	121,369	4540	375	94K	N/A	704	181	6400
frisc	10	1,765,744	3149	2036	164K	443,479	1394	593	3900
misex3	10	N/A	477	493	93K	17,138	477	218	1000
pdc	12	24,552	2865	2810	163K	5839	876	868	4200
s298	9	576,848	1762	1332	192K	11,949	602	294	48,000
s38417	9	489,715	8231	2874	141K	N/A	2248	974	3600
s38584	9	215,098	3422	3283	159K	6605	1946	961	32,000
seq	10	N/A	620	598	96K	5156	573	269	1000
spla	12	14,248	1750	1650	116K	33,618	2802	1387	420
tseng	8	173,711	23,566	385	97K	N/A	N/A	432	890

Speedup in final column is measured at the minimum C achievable by both routers. An “N/A” entry indicates that the router fails to route at the given number of tracks.

be less than half the size of the base switchpoint, suggesting, at most, a 50% switch area penalty for a bit-level network. Using the count net cost heuristic (Fig. 6), we need 10 gates which is likely to be comparable in size to the minimum switch-point implementation. The route assist logic connections are completely local; when switchpoints are wire dominated the area for these additional gates

may be free. As we go to registered and buffered switchpoints (e.g., [33]), the area for the base switch increases, making these additional gates an even smaller marginal cost addition.

The full hardware scheme (Fig. 9) with count net (Fig. 6) and fanout locking requires 30–35 gates. Consequently, this additional logic is likely to be 4× the size

Table 4
Speedup achievable under various hardware-software acceleration schemes at minimum base channels achievable by both routers

Design	Full hardware			Soft allocation		
	Base	Opt.	Parallel	Base	Parallel	Embed
alu4	5000	6700	17,000	130	340	420
apex2	510	720	2200	17	54	54
apex4	28,000	39,000	69,000	950	1600	2900
bigkey	20,000	27,000	94,000	680	2300	2100
clma	1500	2100	6200	39	110	120
des	560	770	1900	18	45	57
diffeq	370	520	1400	12	35	39
dsip	9500	13,000	38,000	320	960	1000
elliptic	160,000	220,000	620,000	5300	15,000	16,000
ex1010	150	210	850	5	21	16
ex5p	6400	8800	14,000	180	300	580
frisc	3900	5300	15,000	100	290	330
misex3	1000	1300	3100	33	75	100
pdc	4200	5800	12,000	140	310	440
s298	48,000	64,000	190,000	1500	4700	4900
s38417	3600	5000	35,000	120	850	380
s38584.1	32,000	45,000	260,000	1000	6400	3300
seq	1000	1500	3100	37	76	110
spla	420	590	1500	14	36	43
tseng	890	1200	3000	27	68	86
geo. Mean	3289	4539	12,851	104	300	327

Speedups measured at same C 's from Table 2. All speedups assume custom hardware. Divide by 10 for an estimate of acceleration for an FPGA-based route engine implementation (Section 10.2).

of the minimal, pass-gate switchpoint. This size is probably untenable for bit-level networks. For multi-bit networks, this area can be amortized across the entire datapath. The inter-compute-page network in SCORE architectures (Section 2.2) for example, uses 4–16b datapaths. Amortized across a 16b datapath, this largest hardware scheme adds about two gate per switchpoint, resulting in only 25–30% area overhead in the worst case.

10.2. FPGA-based spatial router

While it is not viable to embed the full hardware support scheme in a traditional FPGA, it would be feasible to build the hardware for the spatial routing scheme out of FPGA logic. The router logic would be much larger than the FPGA we wanted to route, requiring a board of hundreds of FPGAs to route a single FPGA. This would, for example, allow a logic emulation system (e.g., Quickturn [34], Virtual Wires [3]) to rapidly route each of its component FPGAs.

Implementing the logic equations for a tree switch-point (e.g., Fig. 9) in 4-input Lookup Tables (4-LUTs) will require around 21 4-LUTs. Moreover, we need to consider the additional logic needed at the switch box level such as the random number generator, the parallel prefix circuit (Section 8.2) and various control signals between the switch box and the global route controller. We conservatively estimate that it will require nine more 4-LUTs per switch for the switch-box level logic. As a result, it will require 30 4-LUTs to implement a fast-routing switch.

Conservatively counting a π -switch (two up-links) as two T-switches (one uplink, shown in Fig. 3) and assuming $p = 0.6$ (Table 1), the total number of switchpoints in a design will be:

$$N_{sw} \leq \left(\frac{13}{3}\right) \cdot N_{array} \cdot C \quad (12)$$

See Appendix A for a detailed derivation of Eq. (12). In general, the switch constant, which is $\frac{13}{3}$ here, depends on the particular p value ($0.5 < p < 1.0$). Combining Eq. (12) with the 30 4-LUTs per T-switch equivalent, we see we need about 130 4-LUTs per base channel per leaf tree node. For $C = 13$ (largest value in Table 1), this means the network will have roughly 1700 4-LUTs per leaf tree node. Some additional logic will be required to support the base channel. Since base channels are independent, the logic can easily be sequentialized by base channel, allowing us to save up to a factor of C in LUT count. Many operations (e.g., allocation, victimization, and fanout-free route extensions) need to be performed only on a single base channel tree at a time, so the time increase for this sequentialization will be less than a factor of C .

11. Generalization and future work

The basic path search and allocation strategy introduced here is applicable to any kind of configurable network. As an example, in [19] we present an initial design for hardware-assisted routing for a mesh-based network. The mesh

lacks the convenient, common ancestor switchbox which makes it simple to search from both source and sink. Nonetheless, a single-ended search from the source to the sink is viable in the mesh. A network with purely directional wires (e.g., [26]) would also need a single-ended search and employ variants of the allocation and trace-back techniques introduced for the mesh.

In our route solution, we search all paths from a source to a sink in parallel, and we allow non-overlapping path searches to occur in parallel. We sequentialize all overlapping path searches. In Beneš networks [6], multiBeneš, and multi-fat-trees [2], there are known, parallel algorithms for performing overlapping path searches simultaneously. An interesting area for future work will be to understand if it is possible to expand the kind of hardware-assistance schemes developed here to economically implement overlapping path searches.

12. Conclusions

We have shown that it is possible to design a spatial routing structure that supports arbitrary netlists with fan-out and achieves comparable quality to the history-based Pathfinder route algorithm, the state-of-the-art software routing scheme. Supporting history costs directly in the spatial structure appears prohibitively expensive; however, suitable application of random path selection and route locking appears to be an adequate substitute for Pathfinder's history records. Even when multiple route re-starts are needed to stochastically explore the route space, the parallel, spatial structure can find routes in three to five orders of magnitude fewer cycles than the sequential, software routers. If we must derate this by an order of magnitude to account for the spatial transit time between switch logic when using an FPGA implementation, an FPGA-based spatial router can still be two to four orders of magnitude faster than the software router. This is sufficient to place many routing tasks in the millisecond or sub-millisecond range. Modest schemes that only add hardware for the path search can still be one to three orders of magnitude faster than the pure software approach.

Acknowledgements

This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) (DABT63-C-0048 and ONR N00014-01-0651), the NSF CAREER Program (Grant CCR-0133102), the California MICRO Program, ST Microelectronics, and Xilinx.

Appendix A. Tree switch counting

We can approximate any p value for the HSRA by choosing the sequence of bandwidth preserving (π) switch stages and bandwidth reducing (T) switch stages as described in [12]. Typically this means creating a

repeating sequence of π and T switch stages. If the length of the repeat sequence is l and the number of π switches in the repeat sequence is s , then the Rent exponent (p) is s/l . For $p = 0.6$ as we use in Table 1, we pick $l = 5$ and $s = 3$ and achieve this with the sequence: π, π, T, π, T . There are $\frac{N_{\text{array}}}{2}$ tree switchboxes at the lowest level of the tree, each of which hold C π -switches. One level up, we have $\frac{N_{\text{array}}}{4}$ switchboxes with $2 \cdot C$ π -switches. At the second level up we have $\frac{N_{\text{array}}}{8}$ switchboxes which holds $4 \cdot C$ T-switches. Three levels up we have $\frac{N_{\text{array}}}{16}$ switchboxes which each hold $4 \cdot C$ π -switches. Four levels up we have $\frac{N_{\text{array}}}{32}$ switchboxes which each hold $8 \cdot C$ T-switches. After this stage, the pattern repeats. With this we can write an equation for the switches, identify the geometric sum, and compute the bound on switch count.

$$\begin{aligned}
 N_{\text{sw}} &= C \cdot N_{\text{array}} \cdot \left(\frac{1 \times 2}{2} + \frac{2 \times 2}{4} + \frac{4 \times 1}{8} + \frac{4 \times 2}{16} + \frac{8 \times 1}{32} \right. \\
 &\quad \left. + \frac{8 \times 2}{64} + \frac{16 \times 2}{128} + \frac{32 \times 1}{256} + \frac{32 \times 2}{512} + \frac{64 \times 1}{1024} + \dots \right) \\
 &= C \cdot N_{\text{array}} \cdot \left(1 + 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{4} + \frac{1}{4} + \frac{1}{4} + \frac{1}{8} + \frac{1}{8} + \frac{1}{16} + \dots \right) \\
 &= C \cdot N_{\text{array}} \left(1 + 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{4} + \frac{1}{4} \left(1 + 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{4} \right) \right. \\
 &\quad \left. + \frac{1}{16} (\dots) \right) \\
 &= C \cdot N_{\text{array}} \left(\frac{13}{4} + \frac{1}{4} \left(\frac{13}{4} + \frac{1}{4} (\dots) \right) \right) \\
 &\leq C \cdot N_{\text{array}} \left(\frac{13}{4} \times \left(\frac{1}{1 - \frac{1}{4}} \right) \right) \\
 &\leq C \cdot N_{\text{array}} \left(\frac{13}{4} \times \left(\frac{4}{3} \right) \right) \\
 &\leq C \cdot N_{\text{array}} \left(\frac{13}{3} \right) \tag{13}
 \end{aligned}$$

This finally gives us the result quoted in Eq. (12).

Appendix B. Formal definition of route costs

In this section, we formally define the route cost calculations introduced in Section 6.1. The specifications given here are intended for clarity not for implementation efficiency; Section 6.1 describe how we efficiently implement the cost calculations actually employed.

The route graph for the physical network is a graph $G = (V, E)$. The node set V contains both the endpoints, P , and the switches, S .

$$V = P \cup S \tag{14}$$

The logical design placed on the physical network is a graph $DG = (DV, DE)$.

For each switch, $s \in S$, we have the following fields (see Figs. 2 and 3):

Field	Type	Use
pi	Boolean	$t \rightarrow \pi$ -switch with two parents $f \rightarrow T$ -switch with single parent
$L0$	Boolean	Left downlink connected to parent 0
$L1$	Boolean	Left downlink connected to parent 1 ($pi = t$ case)
$R0$	Boolean	Right downlink connected to parent 0
$R1$	Boolean	Right downlink connected to parent 1 ($pi = t$ case)
LR	Boolean	Left downlink connected to Right downlink
L	$v \in V$	Node attached to left downlink
R	$v \in V$	Node attached to right downlink
$P0$	$s \in S$	Switch attached to Uplink 0
$P1$	$s \in S$	Switch attached to Uplink 1 ($pi = t$ case)
$LNet$	$de \in DE$	Design Net using Left downlink
$RNet$	$de \in DE$	Design Net using Right downlink

We define a distance function, $d(s,p)$, which provides the distance from a switch, $s \in S$, to an endpoint $p \in P$ where each edge is counted as distance one:

$$d(s,p) = \begin{cases} 1, & ((s.L = p) \vee (s.R = p)), \\ \infty, & (s.L \in P \text{ and } s.L \neq p) \text{ and } (s.R \in P \text{ and } s.R \neq p) \\ 1 + \min(d(s.L,p), d(s.R,p)), & \text{otherwise.} \end{cases} \quad (15)$$

B.1. Crossover

The set of possible crossover switches, C , can then be defined in terms of this distance. Each switch in C has equal distance to the source and the sink; further the set C is the set of switches with the minimum such distance. That is, $c \in C$ if and only if:

$$d(c,src) = d(c,sink)$$

and

$$\forall_{s \in S} ((d(s,src) \neq d(s,sink)) \vee (d(s,src) \geq d(c,src)))$$

In all the algorithms used, we identify the subset of crossover switches, $L \subset C$, with least cost from source to sink and then select an $l \in L$ randomly. What differentiates the various schemes consider is how we calculate the cost function which defines L .

B.2. CountCongestion

A Pathfinder-like CountCongestion calculates cost as the number of used segments in the path through the switch c between the source to sink.

$$P_CountCongestion(c,src,sink) \equiv CC(c,src,none) + CC(c,sink,none) \quad (16)$$

In all of our hardware schemes, we take the max of the costs on each side of the crossover rather than the sum.

$$CountCongestion(c,src,sink) \equiv \max(CC(c,src,none), CC(c,sink,none)) \quad (17)$$

To determine which downlink to follow, we define:

$$s.left(p) \equiv (d(s.L,p) < d(s.R,p)) \quad (18)$$

$$s.right(p) \equiv (d(s.L,p) > d(s.R,p)) \quad (19)$$

To calculate when a path is free along the required path, we use:

$$s.free(p,ps) \equiv \begin{cases} s.tfree(p,ps), & \overline{s.pi}, \\ s.xpifree(p,ps), & s.pi \wedge (ps = none), \\ s.l0free(p,ps), & s.pi \wedge (ps = s.P0) \wedge s.left(p), \\ s.l1free(p,ps), & s.pi \wedge (ps = s.P1) \wedge s.left(p), \\ s.r0free(p,ps), & s.pi \wedge (ps = s.P0) \wedge s.right(p), \\ s.r1free(p,ps), & s.pi \wedge (ps = s.P1) \wedge s.right(p). \end{cases} \quad (20)$$

Note that ps is the parent switch along which the route enters the switch in question. This is used here as a mechanism for determining whether the relative uplink is the 0 uplink or the 1 uplink when dealing with a π switch. The special case value of none is used at the crossover where the path does not have a parent switch.

$$\begin{aligned} s.tfree &\equiv \overline{s.LR \vee s.L0 \vee s.R0}, \\ s.xpifree &\equiv \overline{s.LR \vee s.L0 \vee s.R0 \vee s.L1 \vee s.R1}, \\ s.l0free &\equiv \overline{s.LR \vee s.L0 \vee s.R0 \vee s.L1}, \\ s.l1free &\equiv \overline{s.LR \vee s.L1 \vee s.R1 \vee s.L0}, \\ s.r0free &\equiv \overline{s.LR \vee s.L0 \vee s.R0 \vee s.R1}, \\ s.r1free &\equiv \overline{s.LR \vee s.L1 \vee s.R1 \vee s.R0}. \end{aligned}$$

Then, we define path costs to crossover switches as:

$$CC(v,p,ps) \equiv \begin{cases} 0, & v \in P, \\ CC(v.L,p,s), & v \notin P \wedge v.free(p,ps) \wedge v.left(p), \\ CC(v.R,p,s), & v \notin P \wedge v.free(p,ps) \wedge v.right(p), \\ 1 + CC(v.L,p,s), & v \notin P \wedge \overline{v.free(p,ps)} \wedge v.left(p), \\ 1 + CC(v.R,p,s), & v \notin P \wedge \overline{v.free(p,ps)} \wedge v.right(p). \end{cases} \quad (21)$$

B.3. CountNet

The CountNet scheme computes the number of nets intersected on the path from a crossover switch to the source and the sink and takes cardinality of the larger of the two sets as the switch cost.

$$CountNet(c,src,sink) \equiv \max(|CN(c,src,\{c.LNet,c.RNet\})|, |CN(c,sink,\{c.LNet,c.RNet\})|) \quad (22)$$

We count the number of nets with which a path will interfere using an interference set, I , to the crossover switches:

$$CN(v, p, I) \equiv \begin{cases} I, & v \in P, \\ CN(v, p, I \cup \{s.LNet\}), & v \notin P \wedge v.left(p), \\ CN(v, p, I \cup \{s.RNet\}), & v \notin P \wedge v.right(p). \end{cases} \quad (23)$$

B.4. CountNet approximation

The CountNet approximation scheme computes the number of conflicting configurations this net would require between a crossover switch and the source and sink; the maximum of the costs is taken as the cost of using the particular crossover switch.

$$\begin{aligned} & \text{CountNetApproximation}(c, src, sink) \\ & \equiv \max(CNA(c, src, none), CNA(c, sink, none)) \end{aligned} \quad (24)$$

We use $s.left$ and $s.right$ as defined above and further define:

$$s.compat(p, ps) \equiv \begin{cases} s.xtcompat(p, ps), & \overline{s.pi} \wedge (ps = none), \\ s.lcompat(p, ps), & \overline{s.pi} \wedge s.left(p), \\ s.rcompat(p, ps), & \overline{s.pi} \wedge s.right(p), \\ s.xpicompat(p, ps), & s.pi \wedge (ps = none), \\ s.l0compat(p, ps), & s.pi \wedge (ps = s.P0) \wedge s.left(p), \\ s.l1compat(p, ps), & s.pi \wedge (ps = s.P1) \wedge s.left(p), \\ s.r0compat(p, ps), & s.pi \wedge (ps = s.P0) \wedge s.right(p), \\ s.r1compat(p, ps), & s.pi \wedge (ps = s.P1) \wedge s.right(p). \end{cases} \quad (25)$$

$$s.xtcompat \equiv \overline{s.L0} \vee \overline{s.R0}$$

$$s.lcompat \equiv \overline{s.LR} \vee \overline{s.R0}$$

$$s.rcompat \equiv \overline{s.LR} \vee \overline{s.L0}$$

$$s.xpicompat \equiv \overline{s.L0} \vee \overline{s.R0} \vee \overline{s.L1} \vee \overline{s.R1}$$

$$s.l0compat \equiv \overline{s.LR} \vee \overline{s.R0} \vee \overline{s.L1}$$

$$s.l1compat \equiv \overline{s.LR} \vee \overline{s.R1} \vee \overline{s.L0}$$

$$s.r0compat \equiv \overline{s.LR} \vee \overline{s.L0} \vee \overline{s.R1}$$

$$s.r1compat \equiv \overline{s.LR} \vee \overline{s.L1} \vee \overline{s.R0}$$

Then, we define path costs to crossover switches as:

$$CNA(v, p, ps) \equiv \begin{cases} 0, & v \in P, \\ CNA(v.L, p, s), & v \notin P \wedge v.compat(p, ps) \wedge v.left(p), \\ CNA(v.R, p, s), & v \notin P \wedge v.compat(p, ps) \wedge v.right(p), \\ 1 + CNA(v.L, p, s), & v \notin P \wedge \overline{v.compat}(p, ps) \wedge v.left(p), \\ 1 + CNA(v.R, p, s), & v \notin P \wedge \overline{v.compat}(p, ps) \wedge v.right(p). \end{cases} \quad (26)$$

References

[1] Vikas Agarwal, M.S. Hrishikesh, Stephen W. Keckler, Doug Burger, Clock rate versus ipc: the end of the road for conventional

microarchitectures, in: Proceedings of the International Symposium on Computer Architecture, 2000, pp. 248–259.

[2] Sanjeev Arora, F. Tom Leighton, Bruce M. Maggs, On-line algorithms for path selection in a nonblocking network, SIAM Journal on Computing 25 (3) (1996) 600–625.

[3] Jonathan Babb, Russell Tessier, Anant Agarwal. Virtual wires: overcoming pin limitations in fpga-based logic emulators, in: Proceedings of the IEEE Workshop on FPGAs for Custom Computing Machines, April 1993, pp. 142–151.

[4] Vaughn Betz, VPR and T-VPack: versatile packing, placement and routing for FPGAs, <<http://www.eecg.toronto.edu/~vaughn/vpr/vpr.html/>>, March 27 1999. Version 4.30.

[5] Vaughn Betz, Jonathan Rose, FPGA place-and-route challenge, <<http://www.eecg.toronto.edu/~vaughn/challenge/challenge.html/>>, 1999.

[6] Hasan Cam, José A.B. Fortes, Work-efficient routing algorithms for rearrangeable symmetrical networks, IEEE Transactions on Parallel and Distributed Systems 10 (7) (1999) 733–741.

[7] Christopher R. Carroll, A smart memory array processor for two layer path finding, in: Proceedings of the Second Caltech Conference on Very Large Scale Integration, January 1981, pp. 165–195.

[8] Eylon Caspi, Michael Chu, Randy Huang, Nicholas Weaver, Joseph Yeh, John Wawrzynek, André DeHon, Stream computations organized for reconfigurable execution (SCORE): introduction and tutorial, <http://www.cs.berkeley.edu/projects/brass/documents/score_tutorial.html/>, short version appears in FPL'2000 (LNCS 1896), 2000.

[9] Pak K. Chan, Martine D.F. Schlag, Acceleration of an FPGA Router, in: Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines, April 1997, pp. 175–181.

[10] Pak K. Chan, Martine D.F. Schlag, New parallelization and convergence results for nc: a negotiation-based fpga router, in: Proceedings of the International Symposium on Field-Programmable Gate Arrays, ACM/SIGDA, February 2000, pp. 165–174.

[11] S. Christey, The Infinite Monkey Protocol Suite (IMPS). RFC 2759, USC/ISI, Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, California, 90291, April 2000.

[12] André DeHon, Balancing Interconnect and Computation in a Reconfigurable Computing Array (or, why you don't really want 100% LUT utilization), in: Proceedings of the International Symposium on Field-Programmable Gate Arrays, February 1999, pp. 69–78.

[13] André DeHon, Rent's rule based switching requirements, in: Proceedings of the System-Level Interconnect Prediction Workshop (SLIP'2001), ACM, March 2001, pp. 197–204.

[14] André DeHon, Randy Huang, John Wawrzynek, Hardware-Assisted Fast Routing, in: Proceedings of the IEEE Symposium on Field-Programmable Custom Computing Machines, April 2002, pp. 205–215.

[15] S. Even, A. Itai, A. Shamir, On the complexity of timetable and multicommodity flow problems, SIAM Journal on Computing 5 (1976) 691–703.

[16] Stephan W. Gehring, Stefan H.-M. Ludwig, Fast integrated tools for circuit design with fpgas, in: Proceedings of the International Symposium on Field-Programmable Gate Arrays, ACM/SIGDA, February 1998, pp. 133–139.

[17] Ronald I. Greenberg, Charles E. Leiserson, Randomness in computation, volume 5 of Advances in Computing Research, chapter Randomized Routing on Fat-Trees. JAI Press, 1988. Earlier version MIT/LCS/TM-307.

[18] Se June Hong, Ravi Nair, Wire-routing machines—new tools for vlsi physical design, Proceedings of the IEEE 71 (1) (1983) 57–65.

[19] Randy Huang, John Wawrzynek, André DeHon, Stochastic, Spatial Routing for Hypergraph, Trees, and Meshes, in: Proceedings of the International Symposium on Field-Programmable Gate Arrays, February 2003, pp. 78–87.

[20] Randy Ren-Fu Huang, Hardware-Assisted Fast Routing for Runtime Reconfigurable Computing, PhD thesis, University of California at Berkeley, 2004.

- [21] Alexander Iosupovici, A class of array architectures for hardware grid routers, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 5 (2) (1986) 245–255.
- [22] Bradley C. Kuszmaul, Dana S. Henry, Cyclic Segmented Parallel Prefix. UltraScalar Memo 1, Yale, November 1998, <<http://ee.yale.edu/papers/usmemo1.ps.gz>>.
- [23] B.S. Landman, R.L. Russo, On pin versus block relationship for partitions of logic circuits, *IEEE Transactions on Computers* 20 (1971) 1469–1479.
- [24] C.Y. Lee, An algorithm for path connectios and its applications, *IRE Transactions on Electronic Computers* EC-10 (1961) 346–365.
- [25] Tom Leighton, Satish Rao, An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms, in: *Proceedings of the 29th Annual Symposium on the Foundations of Computer Science*, 1988, pp. 422–431.
- [26] Guy Lemieux, Edmund Lee, Marvin Tom, Anthony Yu, Directional and single-driver wires in fpga interconnect, in: *Proceedings of the International Conference on Field-Programmable Technology*, December 2004, pp. 41–48.
- [27] Larry McMurchie, Carl Ebling, PathFinder: a negotiation-based performance-driven router for FPGAs, in: *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, ACM, February 1995, pp. 111–117.
- [28] Stylianos Perissakis, Yangsung Joo, Jinhong Ahn, André DeHon, John Wawrzynek, Embedded DRAM for a reconfigurable array, in: *Proceedings of the 1999 Symposium on VLSI Circuits*, June 1999.
- [29] Christian Plessl, Marco Platzner, Instance-specific accelerators for minimum covering, *Journal of Supercomputing* 26 (2) (2003) 109–129.
- [30] Thomas Ryan, Edwin Rogers, An isma lee router accelerator, *IEEE Design and Test of Computers* (1987) 38–45.
- [31] Jordan S. Swarz, Vaughn Betz, Jonathan Rose, A Fast Routability-Driven Router for FPGAs, in: *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, ACM/SIGDA, February 1998, pp. 140–149.
- [32] Russell Tessier, Negotiated A* Routing for FPGAs, in: *Proceedings of the 5th Canadian Workshop on Field Programmable Devices*, June 1998.
- [33] William Tsu, Kip Macy, Atul Joshi, Randy Huang, Norman Walker, Tony Tung, Omid Rowhani, Varghese George, John Wawrzynek, André DeHon, HSRA: High-Speed, Hierarchical Synchronous Reconfigurable Array, in: *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, February 1999, pp. 125–134.
- [34] Joseph Varghese, Michael Butts, Jon Batcheller, An efficient logic emulation system, *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 1 (2) (1993) 171–174.
- [35] Takumi Watanabe, Hitoshi Kitazawa, Yoshi Sugiyama, A parallel adaptable routing algorithm and its implementation on a two-dimensional array processor, *IEEE Transactions on Computer-Aided Design* 6 (2) (1987) 241–250.
- [36] Michael Wrighton, André DeHon, Hardware-assisted simulated annealing with application for fast FPGA placement, in: *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, February 2003, pp. 33–42.
- [37] Peixin Zhong, Margaret Martonosi, Pranav Ashar, Sharad Malik, Using configurable computing to accelerate boolean satisfiability, *IEEE Transactions on Computed-Aided Design for Integrated Circuits and Systems* 18 (6) (1999) 861–868.