

Go to market faster. Improve quality. Increase competitive advantage.

Cadence NanoRoute Ultra router for Linux on the IBM @server Cluster 1350



Highlights

- **Help reduce time to market and increase competitive advantage with market-leading routing speed**
- **Improve quality of silicon through SMART technology that results in a single, optimized route and fewer design iterations**
- **Improve performance by adding processors with near-linear scalability**
- **Take advantage of potentially lower cost of entry and incremental growth for hardware and software by utilizing the IBM@server Cluster 1350**

As your clients, competitors and the market put pressure on you to develop applications faster and get them to market sooner, routing remains an elusive bottleneck. Design changes require re-routing, delaying project completion and leaving your competitive position vulnerable. Nanometer routers must be able to prevent and correct interdependent routing errors and achieve multiple sign-off quality routing optimization objectives simultaneously. NanoRoute™ Ultra, part of the Cadence® Encounter digital IC platform, running on the IBM @server® Cluster 1350, gives you a way to pass through this routing bottleneck by providing a unified routing and multi-dimensional physical optimization solution.

NanoRoute Ultra is an all-purpose router for block-level and chip-level routing, offering complete post-placement-to-GDSII routing and chip finishing and support for 180 nm, 130 nm and 90 nm process technology rules. It offers one-step rapid timing closure and signal integrity prevention, analysis and correction during detailed routing. By providing more efficient, accurate routing, it helps you achieve more accurate designs the first time through. Combining superthreading technology with IBM **@server** Linux® clusters provides near-linear performance improvements and scalability, helping to improve time to market and encouraging lower total cost of ownership.

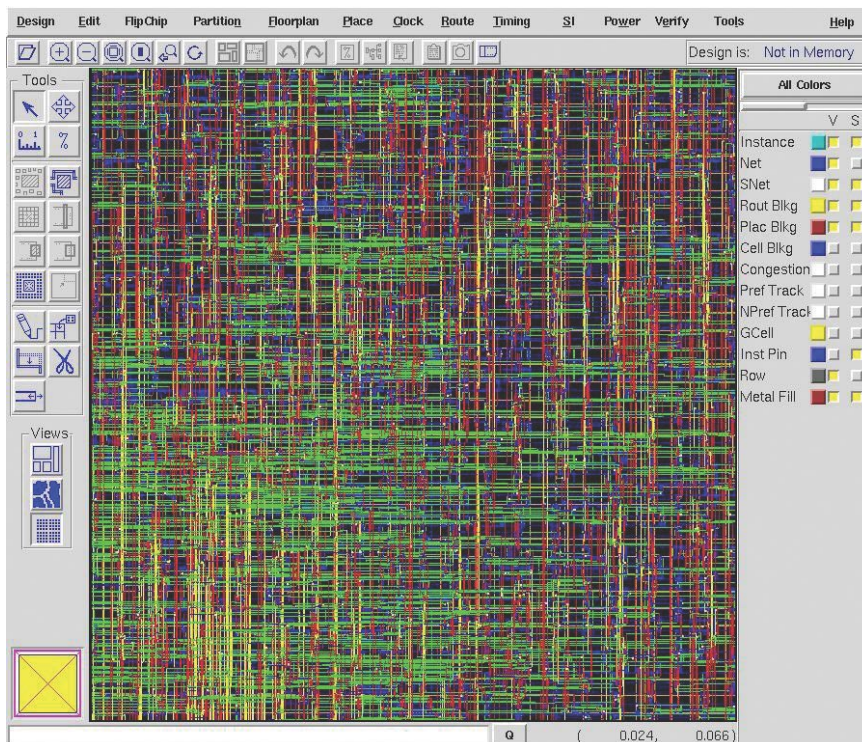
Leverage the fastest routing available and realize performance gains

NanoRoute Ultra from Cadence harnesses the best of distributed and multithreading routing methods to provide an extremely efficient, accurate *superthreading* method. Unlike distributed routing (in which only one CPU is used per machine), superthreading employs every CPU on every machine across the network. The result? Based on internal testing by Cadence, with NanoRoute Ultra, customers can potentially achieve up to 10X routing performance gains in some cases compared to distributed methods¹, in some cases from adding processors while maintaining timing and signal integrity. And with superthreading, since all

CPUs on all machines can be used simultaneously, you can maximize performance while minimizing network traffic to yield high scalability and capacity. Reduce routing time from days to hours or even minutes, bringing applications to market faster and increasing your competitive advantage.

Reduce and predict costs with near-linear scalability

Because superthreading technology enables scaling almost linearly with the number of CPUs available, routing costs can become more stable and predictable, and you can realize improvements in cost-to-performance ratios. For example, in a routing scenario with two CPUs available, a chip designer should see up to a 2X improvement in routing performance; with four CPUs, almost 4X improvement, and so on. This linear scalability lets designers predict the time and cost of routing and help lower the cost versus the performance as more CPUs are added to the network. These advantages make superthreading ideal for large design operations as well as smaller companies who need simple, accessible and affordable routing technology that optimizes current resources.



Magnified view of nanometer routing connecting 60M+ transistors on a digital integrated circuit (IC). NanoRoute Superthreading completes routing in under 2 hours—instead of 38 hours by alternative technologies.

Improve quality of silicon through SMART technology

NanoRoute Ultra is designed to help you run the most efficient, accurate and optimized single route possible, minimizing design iterations and necessary re-routing. This is achieved through concurrent signal integrity, manufacturing-aware, routability and timing-based (SMART) routing technology. Cadence developed this differentiating technology with a simple market reality in mind: that the complexity of today's nanometer designs requires these issues be handled together. For example, a fix on a signal integrity issue will require a change in timing, and vice versa. With SMART technology, you can heuristically manage these optimization objectives, resulting in fewer design iterations, less re-routing and improved quality of results.

Realize the high performance and low cost of Linux clusters on IBM @server systems

The IBM @server Cluster 1350 offers performance and scalability that complement the near-linear performance achieved by NanoRoute Ultra. In a showcase at the Design Automation Conference, Cadence ran NanoRoute Ultra on a Cluster 1350. Using 34 CPUs, NanoRoute Ultra completed the route on a 14-million-gate design in less than two hours. (Running on a competitive architecture, route completion took 38 hours.) The Cluster 1350 offers high performance and availability, which are imperative to chip and application design. Linux clusters also provide a low cost of entry and low total cost of ownership. You can see performance improvements with smaller clusters immediately and scale both hardware and software incrementally as your needs grow. IBM can provide a complete routing

package with NanoRoute Ultra, offering servers along with proven Cluster Systems Management (CSM) software, optional storage and high-speed interconnects to help ease installation and support of your Cluster 1350. For added flexibility, you have the option of running the Cluster 1350 on the architecture that best suits your organization — whether it is Intel®, AMD- or PowerPC-based.

For more information

NanoRoute Ultra on Linux clusters on IBM @server Cluster 1350 offers superior routing, along with improved quality of results and incremental scalability, and enables lower, more predictable design costs.

To learn more about NanoRoute Ultra, visit:
www.cadence.com

To learn more about the IBM @server Cluster 1350, visit:

ibm.com/servers/eserver/clusters/hardware/linux.html

“NanoRoute Superthreading takes routing performance to a whole new level and significantly reduces our design time without sacrificing quality.”

– Jim Wang, Director Design Development, Faraday



© Copyright IBM Corporation 2004

IBM Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States of America
November 2004
All Rights Reserved

¹ In a Cadence performance test, a 14-million-gate design took more than 24 hours to route with one CPU, but finished in less than two hours with superthreading on a Linux network of seven dual processors using all 14 CPUs simultaneously. On a 6-million-gate design, superthreading completed routing in 34 minutes on a network of 20 CPUs, versus more than seven hours on a single CPU.

The examples given in this paper are hypothetical examples of how a customer can use the products described herein and examples of potential cost or efficiency savings are not based on any actual case study. There is no guarantee of comparable results. Many factors determine the sizing requirements and performance of a systems architecture. IBM assumes no liability for the methodology used for determining the configurations recommended in this document nor for the results it provides. Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

IBM reserves the right to change specifications or other product information without prior notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in other countries. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore this statement may not apply to you.

The material included in this document regarding third parties is based on information obtained from such parties. No effort has been made to independently verify the accuracy of the information. This document does not constitute an expressed or implied recommendation or endorsement by IBM of any third-party product or service.

Visit ibm.com/pc/safecomputing periodically for the latest information on safe and effective computing.

Warranty information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, N.C. 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo and @server are trademarks of IBM Corporation in the United States, other countries, or both. For a list of additional IBM trademarks visit ibm.com/legal/copytrade.shtml.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

