Case study

Six competitors, one schedule

How a six-company consortium developed the 20nm process node as a single program — and took three years out of the development schedule

Engagement series.

How lateralworks formed one empowered core team out of IBM, GlobalFoundries, Samsung, STMicroelectronics, Infineon, and the former Chartered Semiconductor — companies that competed everywhere else — and ran the invention of a new process node on a schedule.

Prepared by

lateralworks
FTTM engagement

Sector

Advanced-node semiconductors
Joint development consortium

Online

lateralworks.com
Case study series

Table of contents

Six competitors, one schedule

Core thesis. Invention does not have to arrive on its own clock. When the learning cycles that produce it are named, scheduled, resourced, and run concurrently — and when one person owns each decision — a research program behaves like a delivery program. That is managed innovation: scheduled invention. On 20nm it was worth three years.

Overview

Abstract

In 2010, six major semiconductor companies were trying to invent the same technology at the same time, in the same building, without becoming one team. IBM, GlobalFoundries, Samsung, STMicroelectronics, Infineon, and the former Chartered Semiconductor (by then GlobalFoundries' Singapore operation) were jointly developing the 20nm process node inside a development alliance hosted at IBM's 300mm fab in East Fishkill, New York [1]. Each partner would take the jointly developed process home, add its own IP, and commercialize it in its own fabs [2]. The consortium's leadership had set a two-year target for a node that the program's own arithmetic priced at more than twice that.

The premise was sound. No partner wanted to carry the cost of a leading-edge node alone: process R&D; for the 20nm generation ran to an estimated $2.1–3 billion before a single production fab was built [3]. The practice was harder. Six competitors shared one pilot line, one pool of engineers on rotating assignments, and no single point of decision. Every choice was a negotiation, every negotiation ended in compromise, and the schedule absorbed the cost [1].

GlobalFoundries, the partner with the most riding on the outcome, engaged lateralworks to act as program integrator. We designed and formed a core team drawn from every partner, put a single "product boss" at the head of it, gave each technical area exactly one leader, and built the program's first integrated macro plan — a plan honest enough to show that the technical goals represented roughly five years of work at the program's current learning rate [1].

Then we went to work on the learning rate. More learning cycles, run concurrently rather than in sequence, each one shorter than the last, refreshed weekly against a schedule wired into the fab's manufacturing execution system — the software that runs the factory floor — so the plan updated itself daily. The initial 20nm process, with its design methods and tools, was developed and transferred to every partner in 24 months, roughly three years ahead of the five-year baseline the macro plan had exposed [1]. lateralworks then managed the transfer of the 20nm process into GlobalFoundries Fab 8 in Malta, New York and Fab 1 in Dresden, and on into the development projects that designed and manufactured products for GlobalFoundries' foundry customers [1].

Engagement summary

The engagement at a glance

Role	Program integrator and planning architect — we designed the team and ran the schedule
Client	GlobalFoundries, inside a six-company joint development consortium
Where	IBM’s 300mm development fab, East Fishkill, New York, with transfers to Fab 8 (Malta, NY) and Fab 1 (Dresden)
The target	A new process node in two years, on a technology nobody had built before
The honest read	Our macro plan showed the technical goals represented about five years of work at the current learning rate
The outcome	Initial 20nm process and design methods transferred to all partners in 24 months — three years out of the baseline schedule
Afterward	lateralworks managed the 20nm transfers into Fab 8 and Fab 1, and the product development programs that followed

Company names in this case study are real; the program is a matter of industry record. The names of individual engineers and managers have been removed from the text and blurred in the schedule artifacts.

01

Context
The 20nm moment

To understand what this program attempted, start with where the industry stood in 2010. For forty years, moving to the next process node had been demanding but routine: shrink the transistor, tune the lithography, ship. 20nm was the node where routine ended.

The end of easy scaling

The light used to print chips had stopped shrinking. Immersion lithography at 193nm, the industry's workhorse, runs out of resolution around an 80nm pitch, and 20nm needed features packed at 64nm [4]. The industry's answer was double patterning: split each critical layer onto two masks, expose them separately, and overlay them with nanometer accuracy. Every critical layer now took two passes through the most expensive tools in the fab. Layouts that could not legally be split in two were banned outright, which meant the design rules had to be co-invented with the process itself. The discipline now called design-technology co-optimization (DTCO) became a formal practice at exactly this node [5].

The transistor changed too. In January 2011 the alliance publicly abandoned the gate-first transistor architecture it had championed at 28nm in favor of gate-last — a change in whether the transistor's metal gate is built before or after the highest-temperature processing steps [6]. Gate-first had cost the alliance dearly: at 28nm it did not yield, and TSMC's gate-last approach won that node decisively [7]. At 20nm there was no choice left; restricted design rules and mandatory strain engineering made gate-last the only viable path [6]. The alliance had weighed the two approaches since 2001. Node-defining decisions normally run on decade timescales — which is worth holding in mind against the two-year target this program was given.

What it meant for the participants

The economics were as unforgiving as the physics. A 20nm-generation fab cost $4–7 billion, the process R&D; to fill it ran another $2.1–3 billion, and a single mask set reached $5–8 million [3]. 20nm was also the first node that could not promise cheaper transistors than its predecessor — the point where scaling stopped paying for itself automatically [8]. Only a few giants could still carry a node alone. For everyone else, sharing the bill was the way to stay in the game.

That is what the consortium was. IBM contributed its research pipeline and the East Fishkill line; each partner contributed engineers and money; and each would carry the jointly developed process home, differentiate it with proprietary IP, and compete with the same companies it had just developed it with [2, 9]. For GlobalFoundries the stakes were immediate: a $7 billion greenfield fab was rising in Malta, New York, and the company's credibility as a leading-edge foundry depended on having a leading-edge process to put in it. For Samsung, STMicroelectronics, Infineon, and the rest, the shared process was the admission ticket to the leading edge at a fraction of the solo cost.

Against this backdrop, the consortium's leadership set a two-year target for 20nm technology development [1]. The arithmetic of process development said otherwise. A new process matures through silicon learning

cycles — design an experiment, build masks, run wafers, measure, adjust, go again — and an advanced line typically needs 8 to 12 turns of learning, each historically consuming months of fab time plus a month of characterization [10]. That is two to three years of learning cycles after the process is defined, before qualification even starts. Inside the program, 20nm on a two-year clock was viewed as close to impossible [1]. It was a true pathfinding technology: nobody, anywhere, had built what the roadmap required.

02

The challenge

Six companies, nobody in charge

The alliance concept was good. The operating reality, when we arrived in the fall of 2010, showed why multi-company programs stall — all four of the classic failure modes at once, documented in our field notes from East Fishkill [1].

Nobody in charge

Engineers from the partner companies worked side by side, but each reported to a functional management chain at their home company, off site and with its own priorities. The local hierarchy doubled and sometimes tripled up leadership, with two or three managers “in a box” at each level so that every partner felt represented. Decision-making was slow and convoluted by design, because the design optimized for partner comfort rather than speed. Resources had to be negotiated from the host fab’s heavily siloed, outsourced operating organization.

Competitors on the outside, teammates on paper

The partners cooperated inside the program and competed everywhere else, and the natural behavior followed: take as much learning as possible home, feed as little as possible back. Engineers rotated in and out of the program the way troops rotate through field assignments, and each rotation walked out the door with unrecorded learning.

Scope grew while the clock ran

The program was already late, and most participants privately knew it. In December 2010 the partners responded by expanding the scope: three flavors of 20nm instead of one, three test vehicles instead of one, and three process design kits — PDKs, the models and rules chip designers need before they can design on a process — each with multiple versions [1]. No new resources. No new dates. Tripling scope without adjusting resources or schedules is how a late program becomes a failed one.

A target wearing a schedule’s clothes

Above all of it sat the two-year date. It traced to no sequence of work, no dependency network, no learning-cycle arithmetic; it traced to what the consortium’s leadership needed to be true. lateralworks draws a hard line here. If a date traces back to work, dependencies, and learning cycles owned by the people doing the work, it is a schedule. If it traces back to a board commitment, it is a target wearing a schedule’s clothes [11]. The program had a target. It did not have a schedule. Our job was to give it one, and then to beat it.

The challenge

The alliance in practice

“No single person was in charge. Every partner had to be satisfied, so every decision was a compromise — and the schedule absorbed the cost.”

lateralworks field notes
East Fishkill, February 2011

03

Program architecture

One team out of six companies

The first intervention was not a schedule. It was a team design.

The core team

We designed and formed a core team with the key technical leads from every partner — small, named, and accountable, not the twenty-five-person committee that accumulates when no company wants to be left out [12]. An FTTM core team is a set of program roles whose members subordinate their functions, and in this case their company affiliations, to a shared outcome [13]. Roles came first, names second, company badges last.

One product boss, one leader per area

Six companies were involved, and each wanted to lead every aspect [1]. The single most consequential act of the engagement was establishing single-point management: one “product boss” leading the program, and exactly one named leader for each technical area. Given the political physics of six competitors, getting that governance accepted was itself a major achievement — and everything else depended on it. A program can survive a wrong decision. It cannot survive a decision nobody is empowered to make.

Shared outcomes with doneness criteria

We facilitated the group toward mutually shared outcomes, anchored on key integration milestones with explicit doneness criteria. Each phase gate — technology defined, functional to alpha specification, functional to beta specification, verified to functional specification — carried a published list of deliverables: SRAM yield thresholds, inline parametrics meeting process-capability targets, reliability demonstrations, PDK releases [1]. Doneness criteria turn “are we there yet?” from an argument into a measurement.

20LPM Program Phase Milestones [21Mar11]

ISDA
International Semiconductor Development Alliance

Phase Milestones	Target	Forecast	Deliverables
Definition (Technology defined- M3)	31Jan11	Done 8-Mar11	•Scope / technology targets defined • Digital lib scale factor (0.63x vs 28LP/28G) • Un-tucked PPG suite performance benchmark (+30**%** vs 28LP, +10% vs 28G) •Technology architecture defined •Project plan defined •**Eval 1 PDK** (last deliverable)
Development (functional to Alpha specification – M4)	16Dec11	**31-Mar12**	•RO & RC within 10% of target •25% 8Mb P081 SRAM Vmax yield •80% inline parms (incl d0/HOL) meet Cpk •Reliability capability demonstration •**Alpha PDK**
Prototyping (functional to Beta specification – M5)	28Sep12	**17-Apr13**	•RO & RC within 5% of target •7% 128Mb HD SRAM Vnom Yield •Logic macro yield TBD (target by 15Mar11) •90% inline parms (incl d0/HOL) meet Cpk •Reliability reproducibility •**Beta PDK**
Qualification (verified to functional specification- M6)	29Mar13	**15-Jun13**	•20% 128Mb HD SRAM Vnom yield •Logic macro yield TBD (target by 15Mar11) •**Technology qualification (L1&CPI))** •**500h SRAM stress**: last deliverable.

>2 weeks from target 0-2 weeks from target Forecast at target

4

Figure 1. Program phase milestones with published deliverables per gate — the consortium's doneness criteria as circulated to all partners, March 2011. lateralworks engagement records [1].

Cooperative behavior is modeled, not mandated

You cannot order six competitors to trust each other. You can put their leads in one room, once a week, in front of one plan that shows every partner's contribution and every interface between them, and let the plan referee. First through joint planning and then through weekly refresh planning cycles, cooperation became the path of least resistance. The execution plans made every hand-off explicit, so a slip in one company's scope surfaced as everyone's problem while it was still cheap to fix [1].

The plan opened with a one-line mission agreed by all partners, and it sat at line 1 of the master schedule for the life of the program: *"To enable 20nm bulk technology to be launched into production in Fishkill and/or partner fabs by July 2013."*

ID	Name	Duration	Start	Finish
0	"ISDA 20nm LPM/LPE Technology development"	963 d	9/20/10	6/16/13
1	*To enable 20nm bulk technology to be launched into production in Fishkill and/or partner fabs by July 2013*	0 d	9/20/10	9/20/10
2	Develop 20nm LPM	963 d	9/20/10	6/16/13
3	Definition: Define LPM program scope and technology	187 d	9/20/10	5/2/11
67	Development: Achieve Alpha maturity (IP readiness)	503 d	9/20/10	3/13/12
836	Prototype: Achieve Beta maturity (for IP validation)	833 d	1/3/11	4/22/13
963	Qualification: Verify to functional spec	963 d	9/20/10	6/16/13
1059	Develop 20nm LPM 64nm Mx BEOL	611 d	5/1/11	12/31/12
1060	Definition: Define 64nm Mx BEOL scope	30 d	6/1/11	6/30/11
1069	Development: Achieve alpha maturity (64nm Mx BEOL)	469 d	5/1/11	8/11/12
1123	Prototype: Achieve beta maturity (64nm Mx BEOL)	154 d	5/16/12	10/16/12
1139	Qualification: 64nm Mx BEOLverified to functional spec	169 d	7/16/12	12/31/12
1149	Develop 20nm LPE	884 d	9/20/10	3/29/13
1150	Definition: Define program scope and technology	189 d	9/20/10	5/4/11
1186	Development: Achieve Alpha maturity (IP readiness)	343 d	12/30/10	12/17/11
1262	Prototype: Achieve Beta maturity (for IP validation)	612 d	1/24/11	9/28/12
1308	Qualification: Verify to functional spec	512 d	11/3/11	3/29/13

Figure 2. The integrated master plan for the ISDA 20nm technology development program — one schedule spanning every partner's scope, with the shared mission statement at line 1. lateralworks engagement records [1].

04

The honest diagnosis

The plan that told the truth

A program that intends to schedule invention needs a plan built from the invention itself — the experiments, the dependencies, the learning cycles — not from the date on a board slide.

Programming the invention

In a September 2010 planning workshop we built the program's first integrated macro plan: top level first, seven summary lines running from "define technology" to product development, then progressively decomposed with the people who owned the work [1].

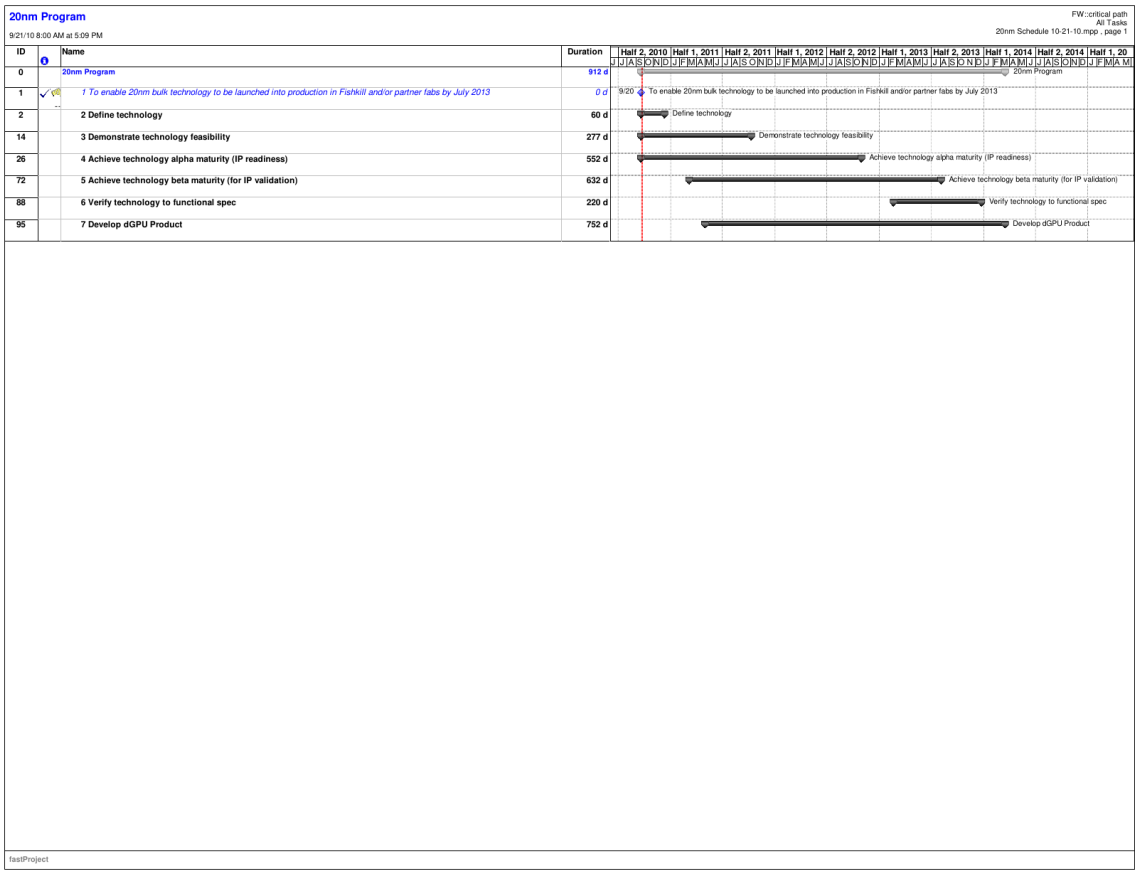

Figure 3. The first cut: the top-level macro plan from the September 2010 planning workshop, seven summary lines spanning definition to product. lateralworks engagement records [1].

The full macro plan grew to more than 1,300 linked activities spanning definition, feasibility, alpha and beta maturity, qualification, and product development, with the critical path computed and visible to everyone. It

was built the FTTM way: invention factored as learning cycles with honest durations, owned by the people doing the work, not as placeholder bars pinned to a board date [11].

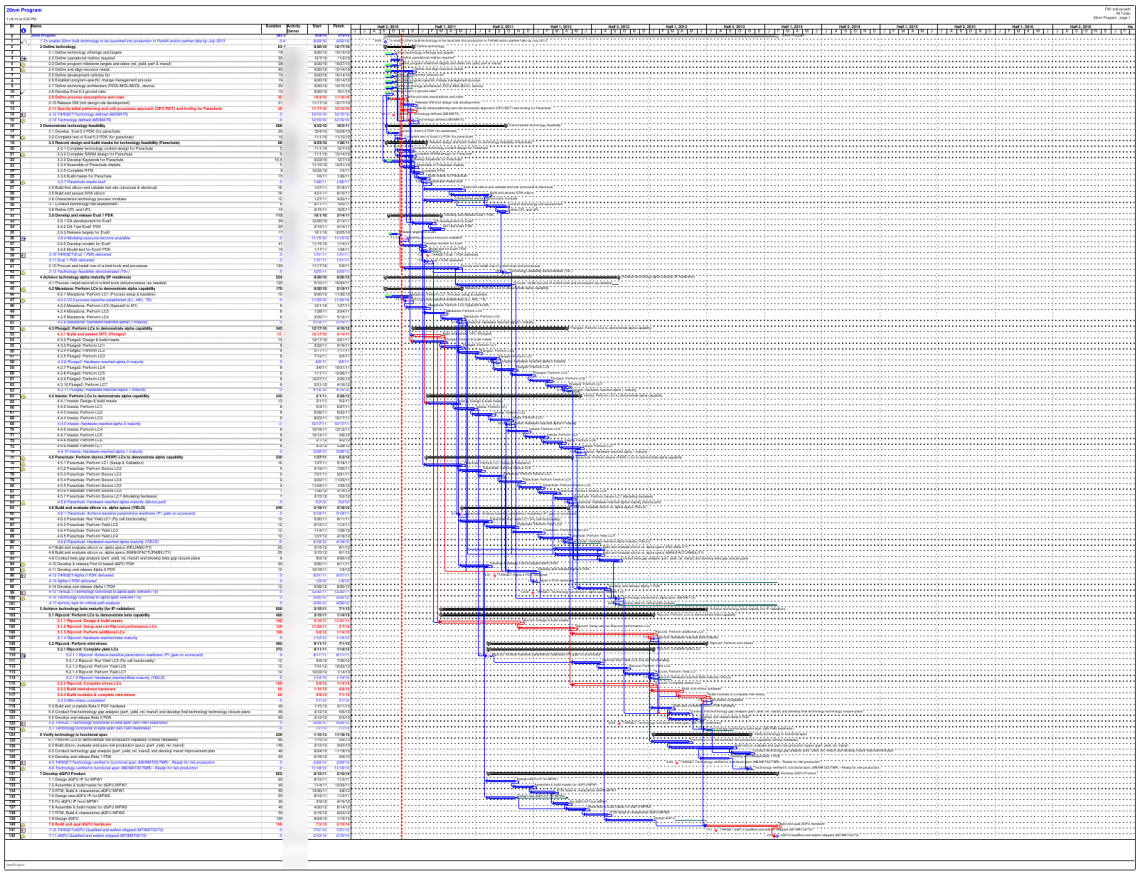

Figure 4. The integrated macro plan, November 2010 — 1,300+ linked activities with the critical path traced in blue and red. Activity-owner initials blurred. lateralworks engagement records [1].

And it told the truth: at the program's current learning rate, the technical goals represented roughly five years of work against the two-year target [1]. Putting that read in front of the team converted an unarguable feeling — "we're late" — into an arguable quantity: we are three years short, and here is where the time is. That is what made acceleration a design problem instead of a morale problem.

The weekly refresh

A plan that is right once is a photograph. A plan that is right every week is an instrument. We ran weekly refresh planning cycles with the core team: every activity owner updated their forecast, the critical path recomputed, and the wigglechart — lateralworks' trend visualization, which plots each milestone's predicted finish date week over week against its target — showed instantly whether the program was pulling in or sliding out.

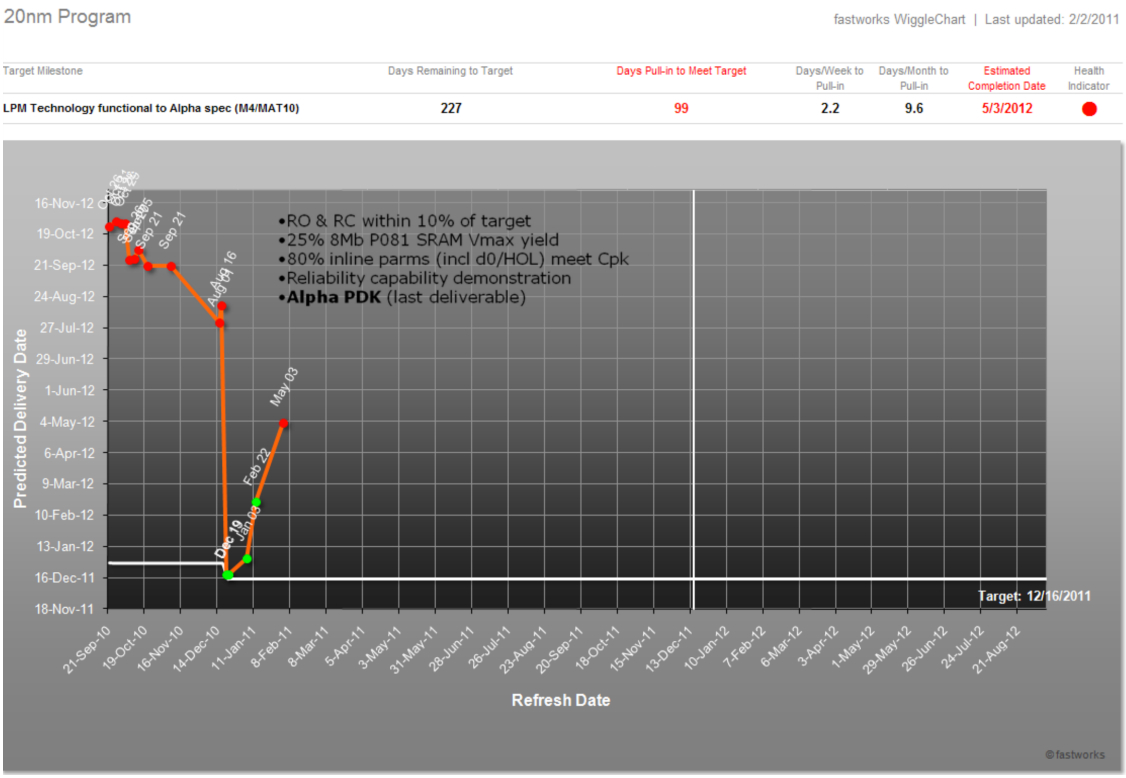

Figure 5. WiggleChart for the alpha-specification milestone, February 2011: the predicted delivery date slid through the fall, snapped back as pull-in actions landed, and carried a computed pull-in requirement of 2.2 days per week to meet target. lateralworks engagement records [1].

The wigglechart made schedule performance a public, weekly fact. Days-per-week of pull-in became a managed number, computed per milestone alongside the gap to target and a health indicator. When a milestone's trend bent the wrong way, the room saw it seven days later, not at the next quarterly review.

20nm Program

fastworks wigglechart | last updated 6/10/2011

Milestone	Target	Estimated Completion	Gap (days)	Completed On	Pull-in (days/week)	Health
LPM Technology verified to functional spec (M6/MAT20/TMR) - Ready for risk production	29-Mar-2013	18-Jul-2013	111		1.2	
LPM Technology functional to Beta spec (M5/TMR readiness)	28-Sep-2012	20-May-2013	234		3.4	
LPM Technology functional to Alpha spec (M4/MAT10)	16-Dec-2011	3-May-2012	139		5.1	

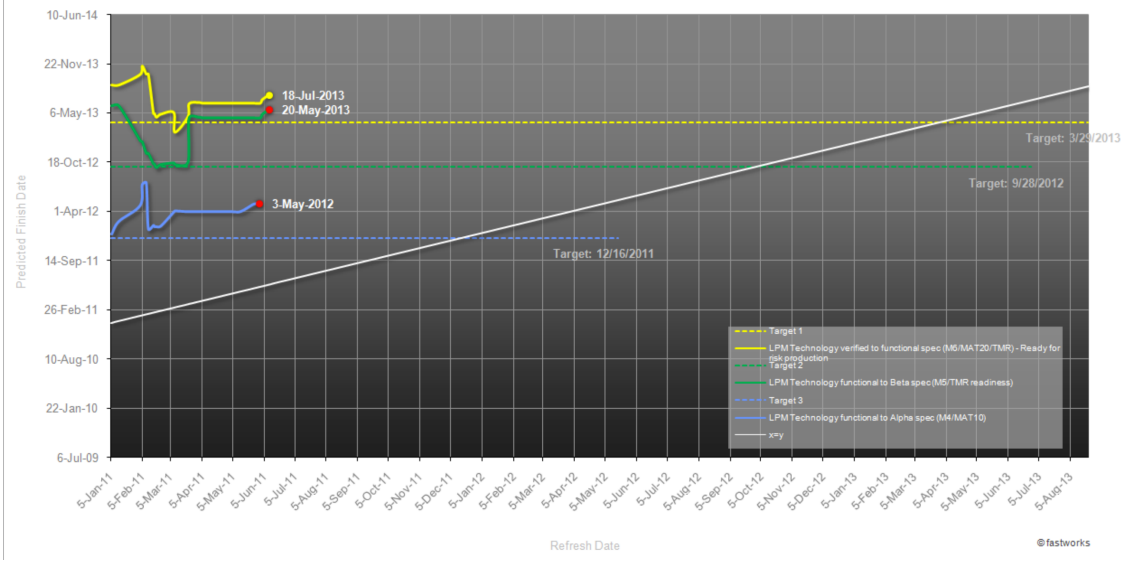

Figure 6. Milestone wigglechart summary, June 2011: three technology milestones tracked weekly, each with gap to target, required pull-in rate, and health status. lateralworks engagement records [1].

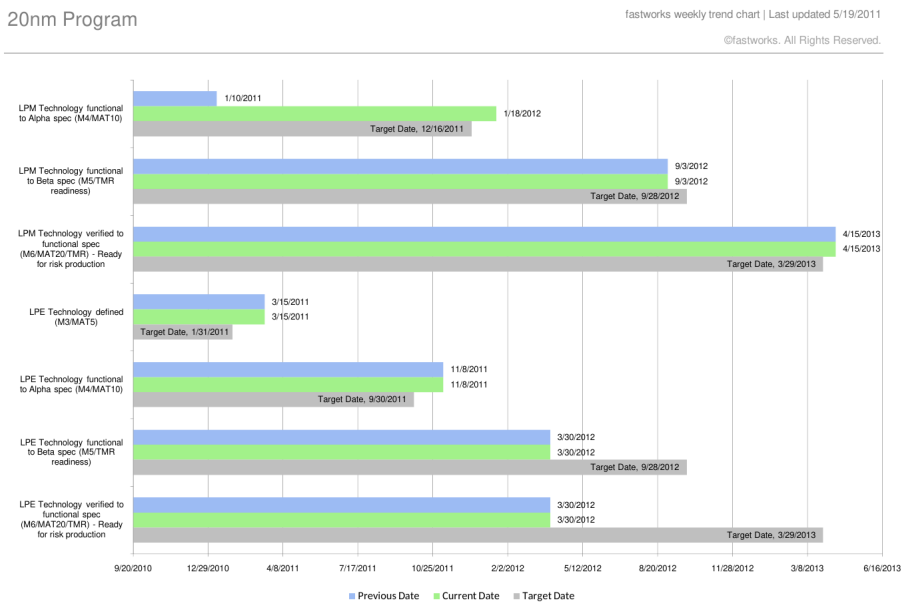

Figure 7. The weekly trend chart circulated to the partners: previous forecast, current forecast, and target date for each technology milestone, May 2011. lateralworks engagement records [1].

Wired into the fab

FTTM schedules stay honest only if updating them is nearly free. We integrated the program schedules with the fab's manufacturing execution system, so lot-level progress — every wafer lot moving through every process step — updated the schedule automatically, every day [1]. Engineers spent their time on engineering. The plan kept itself current.

The honest diagnosis

Target versus schedule

“The macro plan showed five years of work against a two-year target. That gap — not the target — was what we went to work on.”

lateralworks

20nm consortium engagement, 2010

05

Accelerating invention

The learning-cycle engine

Invention in process development has a unit: the silicon learning cycle. Design an experiment, build the masks, run the wafers through the line, measure, learn, go again. Getting an advanced line to production readiness typically takes 8 to 12 turns of learning [10]. The arithmetic is unforgiving — learning cycles multiplied by cycle time *is* the development schedule. Everything else is overhead.

So the acceleration strategy aimed squarely at that arithmetic [1]. **More learning cycles:** more test vehicles and experiments in flight, so more questions were being answered per quarter. **Concurrent learning cycles:** cycles that had been planned in sequence — wait for the answer, then start the next question — were restructured to run in parallel wherever the dependency turned out to be assumed rather than real. **Shorter learning cycles:** priority lots, staged mask releases, and relentless attention to queue time compressed the loop from question to answer.

This is the heart of what lateralworks means by managed innovation — scheduled invention. You cannot schedule a eureka. You can schedule the experiments that produce it: name every learning cycle, give it an owner and an honest duration, put it on the critical path where it belongs, and manage cycle count and cycle time the way a factory manages work in progress. The mystery goes out of the schedule, and the speed goes up.

The infrastructure mattered as much as the intent. The automated lot updater we built kept every experimental lot’s fab progress synchronized with the program schedule daily [1], which meant the learning-cycle engine ran on measured reality rather than on status meetings.

06

Results

What the model delivered

The initial 20nm process, with its design methods and tools, was developed and transferred to each partner in 24 months — against the five-year baseline the macro plan had exposed. Three years came out of the schedule [1].

Where the three years came from

Not from working weekends. The time came out of the structure of the work: one empowered team instead of six delegations, one decision-maker per question, honest plans refreshed weekly, and a learning-cycle engine run flat out. Each partner then took the base technology home and refined it, adding proprietary IP — exactly the operating model the consortium had promised its members [2]. The technology that had looked close to impossible inside the two-year target reached its partners inside a window that kept every one of them in the leading-edge race.

The industry scoreboard

The program tracked itself against TSMC — the largest foundry on earth, running solo with no consortium overhead — milestone by milestone, and by 2011 the consortium's risk-production dates were holding within months of TSMC's [1]. TSMC took 20nm into qualification in 2013 and volume production in 2014 [14]. A jointly developed process, coordinated across six competing companies, held pace on the development clock with the industry leader. That was the entire promise of the alliance model, and for most of the alliance's history the promise had gone unkept.

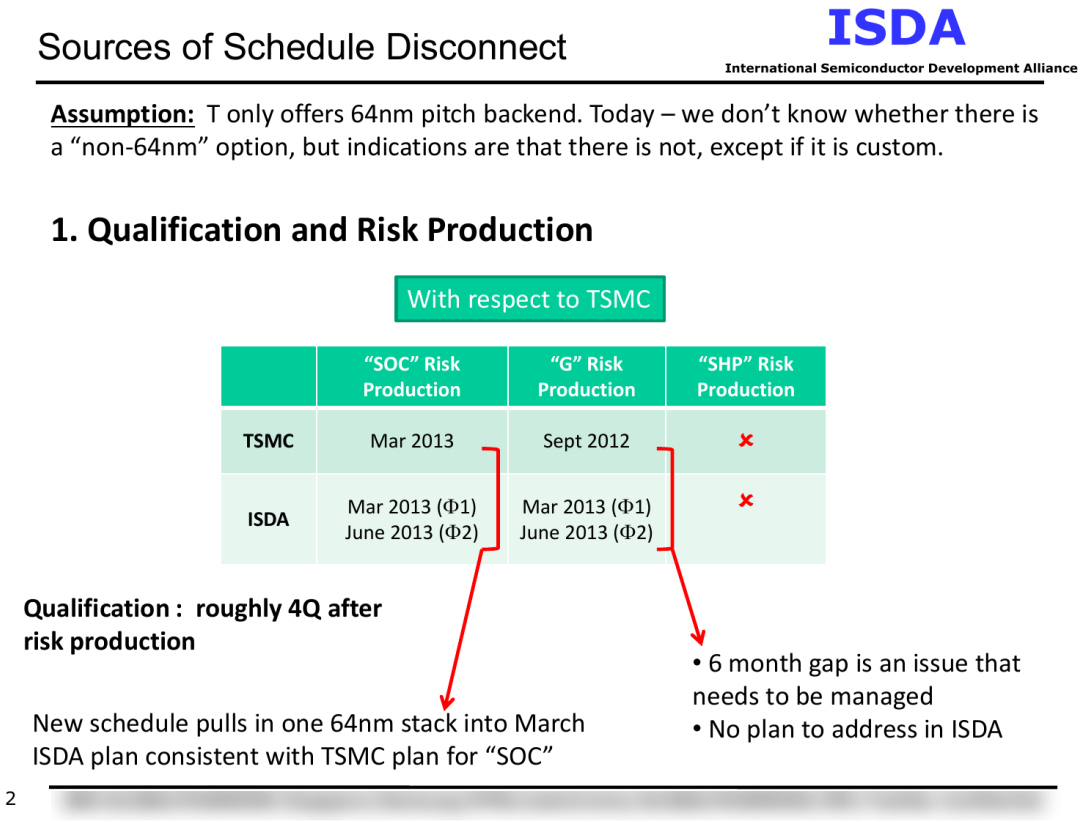

Figure 8. The program's own competitive tracking: consortium risk-production plan versus TSMC, from the schedule dashboard circulated to partners, March 2011. lateralworks engagement records [1].

The transfers

For GlobalFoundries, lateralworks then managed the 20nm process transfer into Fab 8 in Malta, New York and Fab 1 in Dresden — the same integrated-schedule discipline, now spanning technology transfer, tool installation and qualification, and yield learning cycles at the receiving fabs, in one linked model [1]. The engagement continued into the development projects that designed and manufactured products for GlobalFoundries' foundry customers.

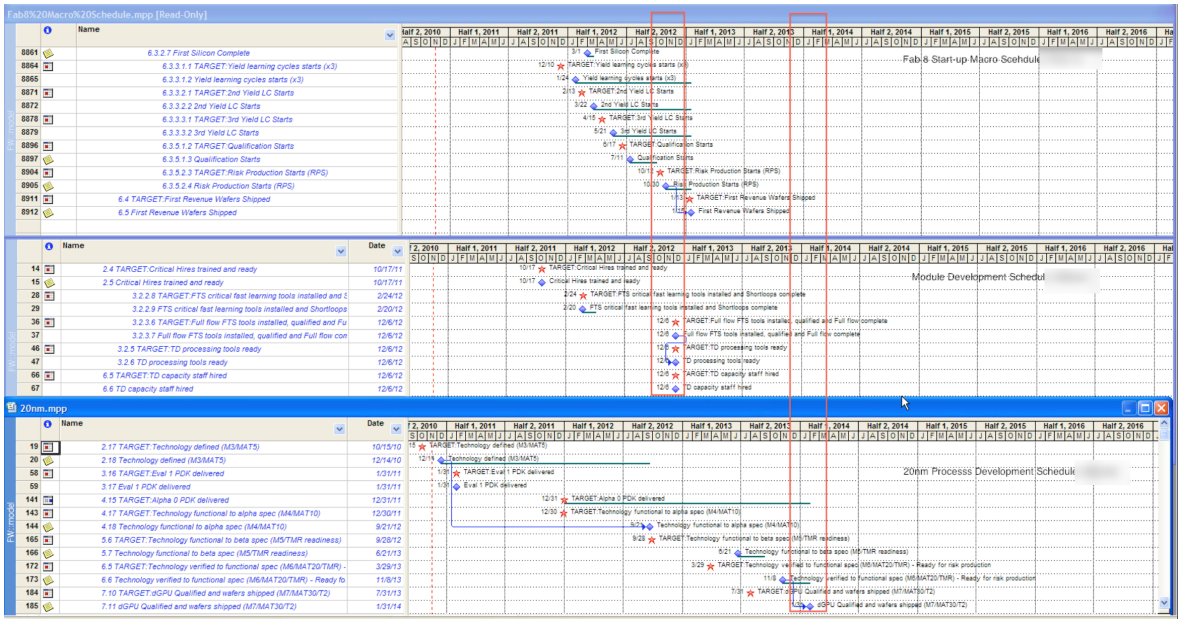

Figure 9. One linked model across three programs: Fab 8 startup macro-schedule, module development, and 20nm process development, integrated so a slip in any one surfaced in all three. Individual names blurred. lateralworks engagement records [1].

What came after

Candor about the node's later history: 20nm planar proved a short commercial generation across the whole industry. FinFET transistors arrived one node later, and in 2014 GlobalFoundries chose to move directly to 14nm [15]. None of that diminishes what the program demonstrated — the development system is what endured. The core-team model, the FTTM schedules wired to the fab, and the wigglechart cadence built here carried directly into the Fab 8 ramp and the transfer programs that followed, where lateralworks' work was later measured at $5 million per day of operational value [16].

07

Best practices

Managed innovation: scheduled invention

This program remains one of lateralworks' clearest demonstrations that FTTM methodology, built on product programs since 1988, applies without dilution to research-grade invention. Seven practices carried it.

1. One empowered core team, drawn from all parties. Roles, not functions: a small named team whose members subordinate company affiliation to a shared outcome [13].

2. Single-point decision governance. One product boss; one leader per technical area. In multi-company programs this is the hardest practice to install and the one that pays for all the others.

3. Mutually shared outcomes with doneness criteria. Integration milestones defined by published deliverables, not by opinion.

4. Honest macro planning. Factor the invention into learning cycles with real durations, owned by the people doing the work. Five years of work does not become two by wishing — it becomes two by redesign [11].

5. Weekly refresh cycles. Joint planning first, then a weekly cadence where forecasts update, the critical path recomputes, and the wigglechart makes the trend public.

6. Learning cycles as the unit of schedule. More cycles, concurrent cycles, shorter cycles. Manage cycle count and cycle time, and the schedule manages itself.

7. Automation at the source of truth. Schedules wired to the fab's manufacturing execution system update daily without human effort. The plan stays alive because staying alive costs nothing.

The through-line is the core thesis: invention scheduled is invention accelerated. Sixteen years on, as the industry again pools development across companies — consortia, joint ventures, and multi-party chiplet programs — the lesson holds unchanged: the constraint is rarely the physics. It is whether many companies can behave as one team. The practices above, and the engagements that proved them, are documented across the lateralworks case study series and methodology library at lateralworks.com.

Sources

References

[1] lateralworks. Engagement records, 20nm consortium program, 2010–2013: integrated macro plans, wigglecharts, weekly trend charts, planning-workshop artifacts, schedule dashboards, and field notes, East Fishkill, NY.

[2] Chartered Semiconductor Manufacturing. Annual report excerpt on joint development of advanced nodes at IBM's East Fishkill facility, SEC Form 6-K, fiscal year 2008. https://www.sec.gov/Archives/edgar/data/0001095270/000114554908000605/u93244exv99w1.htm

[3] McLellan, P. "Cadence at 20nm." SemiWiki, 2012. IBS estimates: $4–7B for 20nm-generation fabs, $2.1–3B process R&D; per node, $5–8M mask sets. https://semiwiki.com/eda/cadence/1591-cadence-at-20nm/

[4] Mishra, P., et al. "Challenges of sub-20nm CMOS." arXiv:1509.00885. Double-patterning requirement below ~80nm pitch; restricted design rules. https://arxiv.org/pdf/1509.00885

[5] Liebmann, L., et al. "The past, present and future of design-technology co-optimization." IEEE CICC, 2013.

[6] James, D. "Common Platform Goes Gate-Last — at Last!" Semiconductor Digest, January 2011. https://www.semiconductor-digest.com/common-platform-goes-gate-last-at-last/

[7] "The Intel Common Platform Foundry Alliance." SemiWiki. On gate-first yield at 28nm and TSMC's gate-last win. https://semiwiki.com/semiconductor-manufacturers/intel/351490-the-intel-common-platform-foundry-alliance/

[8] "High cost per wafer, long design cycles may delay 20nm and beyond." Semiconductor Digest. https://www.semiconductor-digest.com/high-cost-per-wafer-long-design-cycles-may-delay-20nm-and-beyond/

[9] Infineon Technologies. "Infineon, IBM, Chartered, Samsung and Freescale to jointly develop 32nm technology." Press release, May 2007. https://www.infineon.com/cms/en/about-infineon/press/press-releases/2007/INFXX200705-064.html

[10] Gambino, K. "Semiconductor Process Development: Finding a Faster Way to Profitability." SemiEngineering / Coventor, 2017. Silicon learning cycles; 8–12 turns to production readiness. https://semiengineering.com/semiconductor-process-development-finding-a-faster-way-to-profitability/

[11] lateralworks. "Right product, right time." Whitepaper. Targets versus schedules; factoring invention as learning cycles. lateralworks.com/papers

[12] Wheelwright, S. C., and Clark, K. B. Revolutionizing Product Development. Free Press, 1992. Heavyweight team structure.

[13] lateralworks. "The FTTM core team." Whitepaper. Roles, not functions; the heavyweight integrator. lateralworks.com/papers

[14] TSMC. Annual Report 2013: 20nm system-on-chip process qualification and entry into production. https://investor.tsmc.com/static/annualReports/2013/english/c_2_2.html

[15] GLOBALFOUNDRIES. "Samsung and GLOBALFOUNDRIES forge strategic collaboration to deliver multi-sourced 14nm FinFET capacity." Press release, April 2014. https://gf.com/gf-press-release/samsung-and-globalfoundries-forge-strategic-collaboration-deliver-multi-sourced/

[16] lateralworks. "Results." Fab 8 program value measurement. https://lateralworks.com