

PARADE OF TRADES WITH TRANSFER BATCHING

Gregory F. Saragih¹ & Iris D. Tommelein²

ABSTRACT

The Parade of Trades game demonstrates how variability that propagates across sequential trades degrades work flow reliability. Prior work pertaining to the use of takt to manage the Parade of Trades explores how planning for handoffs at a fixed beat can accommodate variability. It also shows where waiting occurs and capacity is lost, and how those influence each trade's time on site and the total project duration. This paper builds on prior work related to the Parade of Trades by adding transfer batching to study how handoff sizing between trades relates to the production system's performance. The study uses discrete-event simulation. Alternative model variants represent capacity variability and vary in transfer-batch sizes. Results show how transfer-batch sizing affects production system performance. Understanding these results can inform work structuring efforts, for example when developing a takt plan or when sizing capacity buffers and timing handoffs between trades.

KEYWORDS

Parade of Trades, Takt planning, variability, discrete-event simulation, transfer batch.

INTRODUCTION

The Parade of Trades game (Tommelein et al., 1999) has become a widely-used learning tool for demonstrating how, due to sequential interdependence of trades, variability propagates and degrades work flow reliability. The game highlights an important insight: even when trades have identical average outputs, variability causes unpredictable handoffs.

Subsequent research has extended the 'Classic' Parade of Trades ('Classic Parade' here refers to the game formulation presented in Tommelein et al., 1999) by exploring game variations, some of which focused on achieving predictable handoffs. Construction practitioners' interest in establishing predictable handoffs contributed to the development of the Last Planner[®] System (Ballard & Tommelein 2021) and, more recently, the increasing adoption of takt planning.

To gain insight into the value of establishing predictable handoffs when designing a production system, the Parade of Trades serves as a reference model and controlled experimental setting in which assumptions about handoff variability can be made explicit, thus enabling comparisons across model variants. In this study, we use discrete-event simulation to compare and contrast production system designs, namely the Classic Parade and two model variants that differ from it in the timing and number of units that get passed from one trade to the next (i.e., the so-called transfer-batch size that defines the handoff).

¹ MS Student, Civil and Envir. Engrg. Dept., Researcher, Project Production Systems Laboratory (P2SL), University of California, Berkeley, CA, USA, gregory.saragih@berkeley.edu, orcid.org/0009-0001-4698-3722

² Roy W. Carlson Distinguished Professor, Civil and Envir. Engrg. Dept., Director, Project Production Systems Laboratory (P2SL), University of California, Berkeley, CA, USA, tommelein@berkeley.edu, orcid.org/0000-0002-9941-6596

The following section describes literature related to the Parade of Trades and defines the term transfer batch. We then present the model assumptions and model variants, and their representation for computer simulation. That section is followed by simulation results, discussion, and conclusions.

LITERATURE

The Classic Parade describes a single-line production system in which sequentially dependent trades pass work through handoffs determined by the roll of a die. The game demonstrates how variability in the quantity of work defined by the outcome of a die roll (contributing to the handoff to the next trade) propagates and affects the downstream trades' use of their capacity, their actual production rate, and time on site, as well as the inventory between trades and the project duration (Tommelein et al., 1999). The variability in the work that gets done by a trade (called within-task variability by Brodetskaia & Sacks, 2007) or process capability (e.g., Tommelein 2020) is captured by the range and likelihood of the die's outcomes. The amplification of variability through sequential trades may be viewed from the chaos perspective on construction (e.g., Bertelsen & Koskela, 2003; Fernandez-Solis, 2013), in which outcomes depend on the system's initial conditions and randomness causes a small deviation from the plan to propagate into large downstream delays. Transfer-batch sizing affects this propagation.

Transfer-batch sizing specifies how much work must be completed before the next trade can start (e.g., one trade must hang some amount of drywall before the follow-on trade can paint it). In flow systems, adherence to a transfer-batch size affects work release timing, inventory between stations, and downstream idle time (Hopp & Spearman, 2008). Smaller transfer batches require more frequent handoffs but enable earlier downstream starts, whereas larger transfer batches reduce handoff frequency at the cost of delayed downstream starts and increased inventory accumulation. With variability, waiting increases as upstream- and downstream work get out of sync, especially when the intention is to keep resource utilization high (Howell et al., 2001; Brodetskaia & Sacks, 2007).

In location-based planning such as takt planning, the definition of zones sets the work package size and therefore functions as a batching decision (Frandsen et al., 2015; Tommelein & Emdanat, 2022). Trades may hand off to the next trade an entire zone of work to be completed within the takt, or perhaps a subset of a zone before full zone completion, which in effect reduces the transfer-batch size (Dlouhy et al., 2019). Thus, zone sizing is a work structuring decision with time and cost implications for the project (Tommelein, 2024).

Buffers provide mechanisms to manage variability in handoffs, albeit at a cost in time, inventory, or capacity (Hopp & Spearman, 2008). In construction, buffering takes the form of providing extra capacity, including time allowances in the schedule, or building up inventory between trades (Dlouhy et al., 2019, Poshdar et al., 2018). Inventory buffers in repetitive construction have been sized to protect downstream work from upstream variability (González et al., 2006) or across zones (Horman et al., 2003). The use of capacity buffers is less explored in construction research compared to the use of time- and inventory buffers (Horman, 2001) but it is highly relevant to takt production systems (Frandsen et al., 2015).

In the Last Planner® System, workable backlog tasks allow trades to work when their planned task is unavailable; they buffer against capacity loss and time loss (Ballard & Tommelein, 2021). When upstream variability disrupts a handoff in terms of timing or quantity, or when remaining work does not require all of a trade's available capacity, that capacity goes unexercised unless workable backlog is available, providing tasks the trade can work on instead.

Buffers come in different kinds as the coordination of specialists requires aligning prerequisites at interfaces, including information, materials, and access, that otherwise translate into downstream waiting (Tommelein & Ballard, 1997). The transfer-batch size defines how much work a trade must complete before handing it off to the follow-on trade, and therefore it

determines how much variability accumulates at each handoff and, accordingly, how much buffering the handoff requires. Transfer-batch sizing can therefore be understood as a design choice for managing variability, reflecting Grote's (2004) distinction between minimizing uncertainty vs. using mechanisms to cope with uncertainty in production system design.

The Parade of Trades game provides a platform to explore production system concepts through simulation. For example, Senior (2011) expands the experimental and measurement features of the Parade by developing an Excel-based Monte Carlo simulation tool, and Han and Park (2011) test reactive managerial actions (e.g., overtime and overmanning) to contrast the firefighting approach with proactive production system design. Tommelein (2020) imposes a planned beat on the Parade logic, and explores the provision of capacity buffers as an upfront—rather than an *ad hoc* decision to manage variability, thereby linking the game (here referred to as the 'Taktet Parade') to takt planning.

Discrete-event simulation, e.g., using EZStrobe (Martinez, 2001), enables systematic exploration of production system design alternatives by allowing assumptions to be stated explicitly and held constant across model variants while outcomes are measured. EZStrobe served as a discovery tool to investigate production system design questions in the Taktet Parade, providing a controlled experimental setting to explore hypotheses and compare model variants (e.g., Rybkowski et al., 2008). Here, we use EZStrobe to study the sizing of transfer batches as a production system design choice in a system subject to variability. The research question is: How does transfer-batch size affect the Parade of Trades?

SIMULATION MODEL

MODEL PARAMETERS AND ASSUMPTIONS

This study presents variants of the Taktet Parade (Tommelein 2020) by imposing transfer batches of a certain size as the handoff from one trade to the next. It addresses the question how zone count Z affects the losses that become visible when using transfer batches, but it does not include standby (or swing) capacity as a buffer. As in the Classic Parade, the new variants represent five trades (A, B, C, D, E) working in sequence to pass all units of work from start to finish, with the capacity of each trade determined by the outcome of rolling a die.

The model variants are parameterized as follows:

- U is the total number of units to pass (e.g., 100)
- Z is the number of transfer batches to pass all units U
- L is the lowest die number, and H is the highest die number (e.g., a L/H die may be a 3/7 die. Deterministic models, i.e., models using the all-five die ($L = H = 5$) without randomness, can provide baseline data for the study.)
- $E = (L+H)/2$ is the expected output per die roll
- $B = U/Z$ is the transfer-batch size, and
- $N = B/E$ is the expected number of rolls needed to produce the transfer batch

Die rolls are independent and identically distributed (IID), meaning that each roll is independent of prior rolls and all trades use the same die.

The transfer-batch size B is a fixed number, chosen by design, that describes how many units a trade must complete before handing off these units to the follow-on trade. Fixed transfer batches require trades to wait until exactly B units are available as input to work on. This waiting will extend the total duration of the Parade. However, transfer batches make handoffs predictable, since each trade knows and can plan for the size of the input they will receive in each handoff. Assuming the handoff is to also occur at a regular time interval, the project can be planned following the beat.

Once a trade has mobilized on site, it may be unable to exercise its capacity for two reasons: (1) the upstream trade has not delivered its handoff that the trade needs as input, or (2) the trade's capacity exceeds the units needed to complete the transfer batch. In the first case, the downstream trade waits for input; the model tracks this as starving. In the second case, the excess capacity will go unused; the model tracks this excess as lost capacity.

For comparison, note that the Classic Parade enforces no transfer batches (no fixed B), i.e., there is a handoff at the end of each roll. When the trade receives as input fewer units than its capacity, due to upstream variability, it will not be able to exercise all of its capacity. According to the aforementioned definitions of starving and lost capacity, these units are tracked as lost capacity (in contrast, Tommelein et al. 1999 used only the term starving).

TRANSFER BATCH CONFIGURATIONS

This study considers two model variants:

$Z = 20$ batches ($B = 5$ units per transfer batch): With $E = 5$, each transfer batch requires $N = B/E = 5/5 = 1$ roll on average, creating 20 transfer batches of 5 units each per trade.

$Z = 10$ batches ($B = 10$ units per transfer batch): With $E = 5$, each transfer batch requires $N = 10/5 = 2$ rolls on average, creating 10 transfer batches of 10 units each per trade.

When transfer batches are passed through the Parade, the structure matches that of a takt plan. In a takt plan, Z would refer to the number of units a trade must complete in one zone, specifically a zone of size $B = U/Z$ units. If every handoff is guaranteed to occur on the takt, the total duration (D) of the deterministic Parade therefore follows the formula $D = (S + Z - 1) \times T(Z)$. With $E = 5$ for all five trades ($S = 5$) to complete $U = 100$:

For $Z = 20$: $D = (S + Z - 1) \times T(Z) = (5 + 20 - 1) \times T(Z) = 24 \times T(Z)$, where $T(Z)$ is the time allotted to one die roll.

For $Z = 10$: $D = (5 + 10 - 1) \times T(Z) = 14 \times T(Z)$, where $T(Z)$ now spans the time allotted to two die rolls.

The corresponding two takt plans are illustrated in Figure 1. Of course, variability in die outcomes will cause the actual duration to deviate from these deterministic values. Furthermore, when a transfer batch of a fixed size is imposed, trades likely will starve as well as lose capacity, as is explained next.

HOW STARVING AFFECTS PARADE PERFORMANCE

When transfer-batching is enforced, a Trade Y will starve when its upstream Trade X requires more rolls than the allotted N rolls to complete its transfer batch and consequently causes a time delay in the handoff to Trade Y. For example, if the transfer-batch size is $B = 5$ units ($N = 1$ roll), when Trade X rolls a 3 it will have to roll a second time to reach B , while the follow-on Trade Y waits. Moreover, if $B = 10$ ($N = 2$), when Trade X rolls a 3 and then 3 again (total = 6), it will have to roll at least a third- and possibly a fourth time. If that third roll produces 7 units when only 4 more units are needed to complete the batch, Trade X also accumulates 3 units of lost capacity. Thus, requiring a transfer batch results in both Trade Y starving and Trade X having lost capacity. The delay caused by additional rolls propagates downstream and extends the Parade's total duration.

HOW CAPACITY LOSS AFFECTS PARADE PERFORMANCE

When transfer-batching is enforced, lost capacity occurs when a trade's die outcome results in more units than the current transfer batch requires. For example, when $B = 5$ is required ($N = 1$) and Trade X rolls 3 units, it falls 2 units short of B (insufficient capacity). If Trade X gets a second roll and then rolls 7 units (cumulative = 10 units), it reaches the required B . Accordingly (assuming Trade X was not starving), the transfer batch from Trade X to Trade Y will comprise 5 units and Trade X's remaining 5 units are tracked as lost capacity. In theory, those 5 units might complete another transfer batch, but the model allows only one handoff at a time. When

$B = 10$ is required ($N = 2$), rolling 7 then 7 (total = 14) completes one batch for Trade X and leaves 4 units of lost capacity. Each unit of lost capacity reflects work Trade X could have completed—had there been work to do—but did not. In the Classic Parade, lost capacity arises differently. Trade X can only hand off as many units as upstream has handed off. When upstream variability results in a handoff to Trade X that is less than what Trade X rolled, lost capacity propagates through the Parade.




Figure 1: Parade with 5 trades and transfer batches, akin to the Taktet Parade.

Top: $Z = 20$ ($B = 5$ units per transfer batch); Bottom: $Z = 10$ ($B = 10$ units per transfer batch)

COMPUTER MODEL

The model is implemented in EZStrobe (Martinez, 2001) version 4.20 with Microsoft Visio Professional 2024, which calls the Stroboscope Simulation System (Version 7,25,03,27). Figure 2 shows an excerpt of the model structure focused on Trade X representing the active trade and Trade Y the follow-on trade. Each trade in the Parade has an identical structure in the model (combis and queues are replicated with trade-specific names).

The EZStrobe model has four parts: (1) Mobilization and Demobilization of the trade, (2) Production Control enforcing handoffs of the required transfer-batch size, (3) Tracking of starving and lost capacity, and (4) Handoff between trades. Capacity variability [DieOutcome] is introduced by rolling a 3/7 die. Note that the priorities assigned to combis play an important part in the order of execution of combis. Thanks to the parametrization of the model, other values can easily be assigned to any of these (e.g., other probability distributions can be used in lieu of the distribution of a fair L/H die).

At the project start time (the time of model initialization), one unit of Trade X is waiting in queue [StartX]. Trade X arrives on site when combi [MobilizeX] ends, releasing the trade to queue [TradeX]. The timing of this mobilization depends on the trade's position in the Parade and is planned based on the expected number of rolls needed to receive its first transfer

batch (e.g., Trade A mobilizes at time 0, Trade B at N , Trade C at $2N$, and so on). Trade X demobilizes when queue $[DoneX]$ accumulates U units. At that time, combi $[DemobilizeX]$ starts and pulls out of its queue $[TradeX]$ so that combi $[RollX]$ will never start again.




Figure 2: Excerpt of EZStrobe model detailing Trade X

When the trade is available in $[TradeX]$, the simulation repeatedly loops from rolling the die $[RollX]$ to waiting to roll again $[WaitToRollX]$, then through $[NextRollX]$ back to $[TradeX]$, until the trade demobilizes. The durations of these combis sum to one time unit (e.g., a day, a week, or whatever time unit one roll represents). $[RollX]$ generates an outcome by evaluating the variable $[DieOutcome]$. This die outcome determines the number of units the trade can produce in that time period, reflecting its capacity, and these units flow to $[WorkX]$. Whether this capacity can be exercised depends on the availability of units in the prerequisite $[InputX]$ (which may result in starving) and subsequently by the units thus far accumulated in $[WorkedX]$ (which may result in lost capacity). If $[InputX]$ has insufficient input, $[WorkX]$ cannot start and Trade X remains idle (it is starving). $[CountStarveX]$ tracks this and accumulates these units into $[StarveX]$. Subsequently, any units completed in excess of B are counted as lost capacity, tracked by $[CountLostCapX]$ and accumulated into $[LostCapX]$.

The control mechanism enforces that exactly B units must be completed before releasing work to the follow-on trade. Once $[WorkedX]$ accumulates at least B units, $[TransferReadyX]$

signals batch completion by checking $[CheckX]$ for the expected number of rolls and releasing B control tokens to $[TargetX]$. One by one, units then transfer through $[TransferBatchX]$ to $[TotalWorkedX]$, which holds the number of units B required for the transfer batch. Handoff $[XY]$ takes B units from $[InputX]$ and transfers the completed transfer batch to $[InputY]$ for follow-on Trade Y.

This study tracks four metrics, three of which (total duration of the Parade, time on site per trade, and lost capacity) align with those tracked by Tommelein (2020). **Total duration** is measured as $SimTime$ when all trades have demobilized. For a given Trade X, **Time on site per trade** is calculated as $XEnd - XStart$, where $XStart$ is the timestamp when $[RollX]$ first starts and $XEnd$ is when $[DemobilizeX]$ starts. **Starving** is measured as $[StarveX].TotCount$. **Lost capacity** is measured as $[LostCapX].TotCount$.

SIMULATION RESULTS

Results were obtained by repeating the simulation for 1,000 iterations to capture stochastic variability, and they are summarized as a mean value $\pm$ one standard deviation. Table 1 describes the Classic Parade and two model variants with data for the first metric: the total duration. Table 2 shows the results for each trade with data for the other three metrics: the trade's (1) total duration on site, (2) starving, and (3) lost capacity. Figures 3 through 6 illustrate the results for each of these four metrics.

Table 1: Total Duration (1,000 iterations)

Model Variant	Z	B	N	D	Avg. $\pm$ SD
Classic	20	3 or 7	1	19 to 38	29.1 ± 1.7
M01	20	5	1	24	41.3 ± 1.8
M02	10	10	2	28	38.0 ± 2.5

Table 2: Time on Site, Lost Capacity, and Starving per Trade (1,000 iterations)

Result	Model Variant	Trade A Avg. $\pm$ SD	Trade B Avg. $\pm$ SD	Trade C Avg. $\pm$ SD	Trade D Avg. $\pm$ SD	Trade E Avg. $\pm$ SD
Time on Site	Classic	20.5 ± 1.9	22.3 ± 1.7	23.4 ± 1.6	24.3 ± 1.7	25.1 ± 1.7
	M01	30.0 ± 2.2	32.6 ± 2.0	34.3 ± 1.9	35.8 ± 1.9	37.3 ± 1.8
	M02	23.9 ± 2.2	26.0 ± 2.3	27.4 ± 2.3	28.8 ± 2.4	30.0 ± 2.5
Starving	Classic	N/A	N/A	N/A	N/A	N/A
	M01	0.0 ± 0.0	12.9 ± 10.7	21.8 ± 11.5	29.5 ± 12.6	36.3 ± 13.3
	M02	0.0 ± 0.0	10.8 ± 10.1	18.7 ± 12.3	25.3 ± 13.2	31.8 ± 14.3
Lost Capacity	Classic	2.5 ± 2.0	10.8 ± 8.6	16.8 ± 9.6	22.1 ± 11.1	25.8 ± 10.9
	M01	49.8 ± 6.6	49.8 ± 6.7	49.9 ± 6.8	49.7 ± 6.8	50.0 ± 6.6
	M02	18.9 ± 6.6	18.8 ± 6.7	18.8 ± 6.8	18.8 ± 6.8	18.6 ± 6.6

TOTAL DURATION

For the deterministic duration of the Classic Parade, the higher value of D reflects rolling a 3 each time resulting in 38 time units, whereas the lower value reflects rolling a 7 each time resulting in 19 time units. For M01 and M02, the computation of the value of D resulting in

values 24 and 28 respectively, appeared earlier in this paper and was illustrated by Figure 1. The simulated Classic Parade completed in 29.1 ± 1.7 time units, whereas both transfer-batch variants took more time: M01 in 41.3 ± 1.8 and M02 in 38.0 ± 2.5 .

Note that, although M02 has a longer deterministic duration than M01, its simulated duration is shorter on average (Figure 3). M01's simulated duration extends 72% beyond its D, because with a single roll allotted per transfer batch, die-roll variability creates batches that 50% of the time require additional rolls to reach B. Likewise, M02's simulated duration extends beyond its D as well; however, its extension of 36% is smaller because the two rolls allotted per transfer batch allow for variability pooling. This is one of two reasons why trades spend less time on site in M02 than they do in M01.

TIME ON SITE PER TRADE

In terms of time on site per trade, the Classic Parade's durations range from 20.5 ± 1.9 (Trade A) to 25.1 ± 1.7 (Trade E), M01's from 30.0 ± 2.2 to 37.3 ± 1.8 , and M02's from 23.9 ± 2.2 to 30.0 ± 2.5 (Figure 4). In all three model variants, the time on site increases with a trade's sequential position in the Parade. Standard deviations vary between ± 1.6 and ± 2.5 across trades, suggesting that variability does not increase downstream even as time on site does.

This pattern reflects the cumulative effect of variability propagating through the Parade. Each subsequent trade inherits the upstream trades' delays and therefore needs more time on site. Trade A shows the least delay since all input units U are available upon mobilization, and Trade E shows the most delay as it carries delays from all four upstream trades. M02 shows shorter times on site than M01 across all trades despite the delay resulting from the required two rolls per batch. As noted, this is because these two rolls allotted per transfer batch allow for variability pooling, which was not possible with a single roll per batch.

STARVING PER TRADE

In terms of starving, expressed in units per trade, M01 ranges from 0.0 ± 0.0 (Trade A) to 36.3 ± 13.3 (Trade E), and M02 from 0.0 ± 0.0 (Trade A) to 31.8 ± 14.3 (Trade E) (Figure 6). In both model variants, starving increases with the trade's sequential position, by roughly 6 to 13 units per position. The first trade in the Parade (Trade A) does not starve because U units are available from the start, whereas each subsequent trade must wait for its upstream trade to complete and hand off a transfer batch. This delay propagates through the Parade: Trade B experiences delay from Trade A, Trade C from both A and B, and so on, with Trade E inheriting the full accumulated delay. Because M02 has fewer handoffs than M01, M02 trades suffer less from the propagation of starving. This is the second reason why trades spend less time on site in M02 than in M01.

LOST CAPACITY PER TRADE

In terms of lost capacity, also expressed in units per trade, the Classic Parade's increases from 2.5 ± 2.0 (Trade A) to 25.8 ± 10.9 (Trade E), as explained previously. In contrast, in M01 and M02, lost capacity is relatively uniform across trades: M01's averages about 50 (range 49.7–50.0) with standard deviations of ± 6.6 to ± 6.8 , and M02's averages about 19 (range 18.6–18.9) with standard deviations of ± 6.4 to ± 6.6 (Figure 5). In transfer-batch model variants, each trade receives a known transfer-batch size B, so lost capacity depends on the die-roll variability that causes time delays to transfer a batch, regardless of the trade's position in the Parade. Lost capacity in M02 is lower than M01 because of variability pooling.




Figure 3: Total Duration of the Parade (1,000 iterations)




Figure 4: Time on Site per Trade (1,000 iterations)




Figure 5: Units of Starving per Trade (1,000 iterations)




Figure 6: Units of Lost Capacity per Trade (1,000 iterations)

DISCUSSION

Comparison of the Classic Parade with M01 and M02 shows a trade-off between the total duration and transfer batch size. The Classic Parade completes faster because trades handoffs occur immediately, but the handoff quantity varies each time. When imposing transfer batches, the handoff quantity is fixed at B, but due to variability in die outcomes additional rolls may be needed to reach the required batch size and the resulting time delays will propagate downstream.

M02 completes faster than M01 in simulation despite having a longer deterministic duration. Two reasons explain this: First, with more rolls per batch, there is opportunity for variability pooling. This reduces the frequency of delayed handoffs. Second, fewer transfer batches mean fewer handoffs at which starving can occur, so less starving propagates downstream, which also leads to less delay to the time on site.

RELEVANCE TO CONSTRUCTION PRACTICE AND TAKT PLANNING

Consider a trade working on 5-unit zones where each zone must be completed before handing it off to the follow-on trade on a set beat (e.g., daily or weekly). Some days the trade completes 3 units (slow day), other days 7 units (fast day). When the trade can complete 7 units but the zone needs only 5, the extra capacity the trade has to produce another 2 is lost capacity (the trade had a fast day but 5 units was enough for the handoff). When the trade completes only 3 and needs a second day to reach 5, the follow-on trade that was scheduled to start on the next beat must wait, i.e., the follow-on trade is starving for work. In the Classic Parade this waiting

does not occur, since whatever is produced gets passed along immediately. However, follow-on trades then receive unpredictable quantities.

When designing a takt plan, the planner must choose a beat that accounts for this variability. Planners may set the takt below the expected average output, for instance at 3 or 4 units per day rather than 5, as an *a priori* capacity buffer against starving of the follow-on trade. The unexercised capacity is then lost when production runs faster than planned. Transfer-batch sizing and takt setting therefore carry cost- and schedule implications: starving and lost capacity both incur direct costs, and more frequent handoffs increase coordination effort.

When using a system like the Last Planner®, workable backlog can put unexercised capacity to productive use. Capacity buffers can ensure that every trade meets the takt so that starving does not propagate downstream.

MATHEMATICAL ANALYSIS

Given the configuration of transfer-batch model variants and binary die outcome, the probability of falling short of B in a transfer batch follows a binomial distribution. For M01, the two possible outcomes are rolling 7, which leaves 2 units of lost capacity, or rolling 3, which falls short and needs another roll to reach $B = 5$. The chance of shortfall is 50%. For M02, the four possible outcomes are 3 then 7, 7 then 3, 7 then 7, and 3 then 3, each with a 25% chance. The first two complete $B = 10$ exactly. The third leaves 4 units of lost capacity, and the fourth falls short, so the chance of shortfall is now only 25%, half that of M01. When the fourth case occurs, rolling 3-3-7 (12.5% chance) produces 13 units and thus 3 units of lost capacity. Rolling 3-3-3-3 (6.25%) or 3-3-3-7 (6.25%) leave 2 and 6 units of lost capacity respectively.

However, making Z even smaller (using larger transfer batches and fewer handoffs) does not reduce the chance of shortfall, as the binomial analysis reveals. For $Z = 5$ ($N = 4$, $B = 20$), a shortfall occurs when none (all 3's) or only 1 of the 4 rolls produces a 7, which is possible for 5 sequences out of 16 (31.25% chance, which is lower than M01's 50% but higher than M02's 25%). For $Z = 2$ ($N = 10$, $B = 50$), a shortfall occurs when fewer than 5 of the 10 rolls produce 7 (386/1024 possible sequences) (37.70% chance), that is, the chance of shortfall increases even as Z decreases beyond M02's $Z = 10$. In summary, increasing N from 1 to 2 decreases the chance of shortfall, but increasing it beyond 2 increases the chance of shortfall, rather than further decreasing it. This non-monotonic behavior suggests that an optimal transfer-batch size exists for a given amount of variability of the die (and thus variability in the expected output).

GENERALIZABILITY

Although the transfer-batch model variants in this study use five trades in a linear sequence, the underlying dynamics apply to larger and more complex Parades as well. The study uses a deliberately simplified model with a binary die to isolate the effect of transfer-batch sizing under controlled conditions and to enable the binomial analysis. Real-world production systems exhibit richer variability profiles, and a trade's process capability (i.e., the probability distribution of its output) may not be well known *a priori*, making careful data collection a prerequisite for identifying an appropriate transfer-batch size in practice. Enforcing transfer batches under high variability can amplify the consequences across the Parade. Proper transfer-batch sizing is necessary when managing variability in a Parade.

CONCLUSIONS

This study used discrete-event simulation to explore how transfer batching affects the Parade of Trades. Two transfer-batch model variants were compared to the Classic Parade, tracking total duration as well as time on site, starving, and lost capacity per trade across 1,000 iterations.

The results show that transfer batching makes handoff quantities predictable but is likely to extend the total duration because a trade must complete a transfer batch before handing it off

to the follow-on trade. The binomial analysis reveals that the relationship between transfer-batch size and chance of shortfall is not monotonic, suggesting that an appropriate transfer-batch size depends on the variability profile and therefore the expected output. Transfer-batch sizing must take into account variability in process capability (as modeled by the die roll). Predictable handoffs in timing and quantity are particularly important in the process of takt planning. The parametrized model presented in this study provides a controlled setting in which such relationships can be explored, both as a learning tool and as a research platform.

Future work could explore the provision of buffers such as capacity- or time buffers as variability management strategies, transfer-batch configurations with different numbers of transfer batches and die probability profiles, and work flows beyond a single trade sequence where completion of a transfer batch may be required as input to one or multiple parades.

Relating performance metrics to costs, allowing work carryover between transfer batches or using workable backlog, and validating findings against planning practice would strengthen the practical applicability of these insights. This study invites practitioners and researchers to reflect on how zoning decisions in (takt) planning affects variability during execution, and to explore other work structuring alternatives that build upon the Parade of Trades.

ACKNOWLEDGMENTS

This study was supported by financial contributions from members of the Project Production Systems Laboratory (P2SL) at University of California, Berkeley. We gratefully acknowledge all support received. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of P2SL members.

REFERENCES

- Ballard, G., & Tommelein, I. D. (2021). *2020 Current Process Benchmark for the Last Planner® System of Project Planning and Control*. Technical Report, Project Production Systems Laboratory (P2SL), University of California, Berkeley, California, USA. doi.org/10.34942/P2F593
- Bertelsen, S., & Koskela, L. (2003). Avoiding and Managing Chaos in Projects. *Proc 11th Ann. Conf. Int. Group for Lean Constr.* (IGLC11). Virginia, USA. iglc.net/Papers/Details/229
- Brodetskaia, I., & Sacks, R. (2007). Understanding Flow and Micro-Variability in Construction: Theory and Practice. *Proc. 15th Ann. Conf. Int. Group for Lean Constr.* (IGLC15). Michigan, USA. 488-497. iglc.net/Papers/Details/478
- Dlouhy, J., Binninger, M., & Haghsheno, S. (2019). Buffer Management in Takt Planning – An Overview of Buffers in Takt Systems. *Proc. 27th Ann. Conf. Int. Group Lean Constr.* (IGLC27). Dublin, Ireland. 429-44. doi.org/10.24928/2019/0226
- Fernandez-Solis, J. L. (2013). Building construction: A deterministic non-periodic flow – A case study of chaos theories in tracking production flow. *Architectural Engineering and Design Management*, 9(1), 21-48. doi.org/10.1080/17452007.2012.683671
- Frandsen, A. G., Seppänen, O., & Tommelein, I. D. (2015). Comparison between Location Based Management and Takt Time Planning. *Proc. 23rd Ann. Conf. Int. Group Lean Constr.* (IGLC23). Perth, Australia. 3-12. iglc.net/Papers/Details/1181
- González, V., Alarcón, L. F. & Gazmuri, P. (2006). Design of Work in Process Buffers in Repetitive Building Projects: A Case Study. *Proc. 15th Ann. Conf. Int. Group for Lean Constr.* (IGLC14) Santiago, Chile. 165–176. iglc.net/papers/Details/424
- Grote, G. (2004). Uncertainty management at the core of system design. *Annual Reviews in Control*, 28(2), 267-274. doi.org/10.1016/j.arcontrol.2004.03.001
- Han, S., & Park, M. (2011). Interactive parade game: Impact of managerial reactions to workflow variability. *J. Inf. Technol. Constr.*, 16, 105–118. Retrieved from itcon.org/papers/2011_7.content.01802.pdf

- Hopp, W. J., & Spearman, M. L. (2008). *Factory physics* (3rd ed.). McGraw-Hill/Irwin.
- Horman, M. J. (2001). Modeling the Effects of Lean Capacity Strategies on Project Performance. *Proc. 9th Ann. Conf. Int. Group Lean Constr.* (IGLC9). Singapore. iglc.net/Papers/Details/144
- Horman, M. J., Messner, J. I., Riley, D. R., & Pulaski, M. H. (2003). Using Buffers to Manage Production: A Case Study of the Pentagon Renovation Project. *Proc. 11th Ann. Conf. Int. Group for Lean Constr.* (IGLC11). Virginia, USA. iglc.net/Papers/Details/240
- Howell, G. A., Ballard, G., & Hall, J. (2001). Capacity Utilization and Wait Time: A Primer for Construction. *Proc. 9th Ann. Conf. Inter. Group Lean Constr.* (IGLC9). Singapore. iglc.net/Papers/Details/145
- Howell, G., Laufer, A., & Ballard, G. (1993). Interaction between Subcycles: One key to Improved Methods. *J. Constr. Engin. Manage.*, 119(4), 714-728. [doi.org/10.1061/\(ASCE\)0733-9364\(1993\)119:4\(714\)](https://doi.org/10.1061/(ASCE)0733-9364(1993)119:4(714))
- Martinez, J. C. (2001). EZStrobe: General-purpose simulation system based on activity cycle diagrams. *Proc. 2001 Winter Sim. Conf.* 1556-1564. doi.org/10.1109/WSC.2001.977485
- Poshdar, M., Gonzalez, V., & Kasiviswanathan, B. (2018). Buffer Management in Construction - a New Zealand Study. *Proc. 26th Ann. Conf. Int. Group Lean Constr.* (IGLC26). Chennai, India. 818-828. doi.org/10.24928/2018/0517
- Rybkowski, Z. K., Wong, J., Ballard, G., & Tommelein, I. D. (2008). Using Controlled Experiments to Calibrate Computer Models: The Airplane Game as a Lean Simulation Exercise. *Proc. 16th Ann. Conf. Int. Group Lean Constr.* (IGLC16). Manchester, UK. 309-319. iglc.net/Papers/Details/606
- Senior, B. A. (2011). Enhancing the construction Parade of Trades game. *Proc. 47th ASC Ann. Int. Conf.* Univ. of Nebraska-Lincoln in Omaha, NE.
- Tommelein, I. D. (2020). Taktung the Parade of Trades: Use of Capacity Buffers to Gain Work Flow Reliability. *Proc. 28th Ann. Conf. Int. Group for Lean Constr.* (IGLC28). Berkeley, California, USA. 421-432. doi.org/10.24928/2020/0076
- Tommelein, I. D. (2024). About Time-Cost Trade-Offs in Takt Planning. *Proc. 32nd Ann. Conf. Int. Group Lean Constr.* (IGLC 32). Auckland, New Zealand. 226-237. doi.org/10.24928/2024/0157
- Tommelein, I. D., & Ballard, G. (1997). *Coordinating specialists*. University of California, Berkeley. doi.org/10.34942/P2G59D
- Tommelein, I. D., & Emdanat, S. (2022). Takt Planning: An Enabler for Lean Construction. *Proc. 30th Ann. Conf. Int. Group Lean Constr.* (IGLC30). Edmonton, Alberta, Canada. 866-877. doi.org/10.24928/2022/0198
- Tommelein, I. D., Riley, D. R., & Howell, G. A. (1999). Parade game: Impact of work flow variability on trade performance. *J. Constr. Engin. Manage.*, 125(5), 304-310. [doi.org/10.1061/\(ASCE\)0733-9364\(1999\)125:5\(304\)](https://doi.org/10.1061/(ASCE)0733-9364(1999)125:5(304))